Spectral Simplicity of Apparent Complexity, Part II: Exact Complexities and Complexity Spectra

Paul M. Riechers* and James P. Crutchfield[†] Complexity Sciences Center Department of Physics University of California at Davis One Shields Avenue, Davis, CA 95616 (Dated: January 3, 2018)

The meromorphic functional calculus developed in Part I overcomes the nondiagonalizability of linear operators that arises often in the temporal evolution of complex systems and is generic to the metadynamics of predicting their behavior. Using the resulting spectral decomposition, we derive closed-form expressions for correlation functions, finite-length Shannon entropy-rate approximates, asymptotic entropy rate, excess entropy, transient information, transient and asymptotic state uncertainty, and synchronization information of stochastic processes generated by finite-state hidden Markov models. This introduces analytical tractability to investigating information processing in discrete-event stochastic processes, symbolic dynamics, and chaotic dynamical systems. Comparisons reveal mathematical similarities between complexity measures originally thought to capture distinct informational and computational properties. We also introduce a new kind of spectral analysis via coronal spectrograms and the frequency-dependent spectra of past-future mutual information. We analyze a number of examples to illustrate the methods, emphasizing processes with multivariate dependencies beyond pairwise correlation. An appendix presents spectral decomposition calculations for one example in full detail.

PACS numbers: 02.50.-r 89.70.+c 05.45.Tp 02.50.Ey 02.50.Ga

Keywords: hidden Markov model, entropy rate, excess entropy, predictable information, statistical complexity, projection operator, complex analysis, resolvent, Drazin inverse

0

CONTENTS

1.	Introduction	- 2
	A. Notational review	2
	B. Outline of main results	3
II.	Correlation and Myopic Uncertainty	4
	A. Nonasymptotics	4
	B. Asymptotic correlation	5
	C. Asymptotic entropy rate	6
III.	Accumulated Transients for Diagonalizable	
	Dynamics	6
IV.	Exact Complexities and Complexity Spectra	8
	A. Excess entropy	8
	B. Persistent excess	9
	C. Excess entropy spectrum	9
	D. Synchronization information	10
	E. Power spectra	11
	F. Almost diagonalizable dynamics	11
	G. Markov order versus symmetry collapse	12
V.	Spectral Analysis via Coronal Spectrograms	12
VI.	Examples	14
	A. Golden Mean Processes	14

[†] chaos@ucdavis.edu

T T 4 1 4.

	B. Even Process	15
	C. Golden–Parity Process Family	16
VII.	Predicting Superpairwise Structure	18
VIII.	Conclusion	21
	Acknowledgments	21
А.	Example Analytical Calculations	22
	1. Process and spectra features	22
	2. Observed correlation	23
	3. Predictability	25
	4. Synchronizing to predict optimally	27
	References	28

The prequel laid out a new toolset that allows one to analyze in detail how complex systems store and process information. Here, we use the tools to calculate in closed form almost all complexity measures for processes generated by finite-state hidden Markov models. Helpfully, the tools also give a detailed view of how subprocess components contribute to a process' informational architecture. As an application, we show that the widely-used methods based on Fourier analysis and power spectra fail to capture the structure of even very simple structured pro-

^{*} pmriechers@ucdavis.edu

cesses. We introduce the spectrum of past-future mutual information and show that it allows one to detect such structure.

I. INTRODUCTION

Tracking the evolution of a complex system, a time series of observations often appears quite complicated in the sense of temporal patterns, stochasticity, and behavior that require significant resources to predict. Such complexity arises from many sources. Apparent complexity, even in simple systems, can be induced by practical measurement and analysis issues, such as small sample size, inadequate collection of probes, noisy or systematically distorted measurements, coarse-graining, out-ofclass modeling, nonconvergent inference algorithms, and so on. The effects can either increase or decrease apparent complexity, as they add or discard information, hiding the system of interest from an observer to one degree or another. Assuming perfect observation, complexity can also be inherent in nonlinear stochastic dynamical processes-deterministic chaos, superexponential transients, high state-space dimension, nonergodicity, nonstationarity, and the like. Even in ideal settings, the smallest sufficient set of a system's maximally predictive features is generically uncountable, making approximations unavoidable, in principle [1]. With nothing else said, these facts obviate physical science's most basic goal—prediction—and, without that, they preclude understanding how nature works. How can we make progress?

The prequel, Part I, argued that this is too pessimistic a view. It introduced constructive results that address hidden structure and the challenges associated with predicting complex systems. It showed that questions regarding correlation, predictability, and prediction each require their own analytical structures. Measures of correlation—including autocorrelation functions, power spectra, and Greeen-Kubo transport coefficients-are direct signatures of the transition dynamic of any hidden Markov model (HMM) representation of a process. However, Part I explained that synchronizing to the hidden linear dynamic in systems *induces* a nondiagonalizable metadynamics, even if the dynamics are diagonalizable in their underlying state-space. Hence, questions about predictability—for example: "How much of the future can be predicted from past observations?" and "What is the irreducible randomness of the process?"—and the burdens of actually predicting—"How much memory must be allocated to predict what is predictable?" and "How much time must be invested before an observer is sufficiently synchronized to make good predictions?"-

are instead answered by the nondiagonalizable transition dynamics of so-called mixed state presentations (MSPs) of the process.

Part I gave operator expressions that begin to answer these questions. However, assuming normal and diagonalizable dynamics, so familiar in mathematical physics, simply fails in this setting. Thus, nondiagonalizable dynamics presented an analytical roadblock. Part I reviewed a calculus for functions of nondiagonalizable operators—the recently developed meromorphic functional calculus of Ref. [2]—that directly addresses nondiagonalizability, giving constructive calculational methods and algorithms. Part I also highlighted a spectral weighted directed-graph theory that can give useful shortcuts for determining a process' spectral decomposition.

Part II now goes a step further, providing the first closed-form expressions for many of the complexity measures in wide use. In following through the derivations, we discover a timescale—the *symmetry collapse index*—that indicates the sophistication of finite computational structures in infinite-Markov-order processes. We are also led to the introduction of *complexity spectra* that give a frequency-decomposition of predictable features.

In the following, we consider wide-sense stationary stochastic processes $\dots X_{t-1}X_tX_{t+1}\dots$ with a discrete domain. The random variable X_t can take on values $x \in \mathcal{A}$, which we assume to be a finite set. The interdependencies among the random variables can be interpreted as intrinsic computations carried out by the process under observation. While such stochastic processes are natural in discrete-time signal processing [3, 4] and on discrete lattices [5], they are sometimes also obtained from continuous-time chaotic systems via regular sampling. Moreover, while finite alphabets \mathcal{A} are intrinsic to certain sources, they may also be obtained from continuous phase-spaces, either by measurement precision limitations or by intentional parsimony as in the case of generating partitions used in symbolic dynamics [6–8]. Such processes are typically non-Markovian and can appear very complicated, but we unravel their complexities via analysis of a HMM representation of the process where the HMM may be obtained from first principles or may be inferred by any number of methods (see Ref. [9] and references therein).

A. Notational review

Part II here uses Part I's notation and assumes familiarity with its results.

In Part I, we reviewed relevant background in stochastic processes and their complexities and the hidden Markov models (HMMs) that generate them. Part I delineated several classes of HMMs-Markov chains, unifilar HMMs, and nonunifilar HMMs. It also reviewed their mixed-state presentations (MSPs)—HMM generators of a process that track distributions induced by observation. Related constructions included the mixedfunctional presentations and cryptic-operator presenta-MSPs are key to calculating complexity meations. sures within an information-theoretic framing. Part I then showed how each complexity measure reduces to a linear algebra of an appropriate HMM adapted to the cascading- or accumulating-question genre. It summarized the meromorphic functional calculus and several of its mathematical implications in relation to spectral projection operators.

Recall that a distribution η over the states \boldsymbol{S} of an HMM $\mathcal{M} = (\boldsymbol{S}, \mathcal{A}, \{T^{(x)}\}_{x \in \mathcal{A}}, \mathcal{S}_0 \sim \eta)$ induces a probability distribution over subsequent words $x_{0:L} = x_0 x_1 \dots x_{L-1}$ generated by the process:

$$\Pr_{\mathcal{S}_0 \sim \eta} (X_{0:L} = x_{0:L}) = \langle \eta | T^{(x_0:L)} | \mathbf{1} \rangle$$
$$= \langle \eta | T^{(x_0)} T^{(x_1)} \dots T^{(x_{L-1})} | \mathbf{1} \rangle ,$$

where the labeled transition matrices, with elements $T_{i,j}^{(x)} = \Pr(X_t = x, \mathcal{S}_{t+1} = s_j | \mathcal{S}_t = s_i)$, sum to the rowstochastic transition matrix $T = \sum_{x \in \mathcal{A}} T^{(x)}$, and $|\mathbf{1}\rangle$ is the all-ones vector such that $T |\mathbf{1}\rangle = |\mathbf{1}\rangle$.

Ergodic processes will have a single stationary distribution π , such that $\langle \pi | T = \langle \pi |$. (Non-ergodic processes can still have a specified stationary distribution π over states, but it could no longer be uniquely identified from T alone.) The synchronizing MSP (S-MSP) introduced in Part I is a metadynamic of how distributions over \mathcal{M} are updated by observation, starting in the mixed state π and thus starting in the peaked distribution δ_{π} over the set of mixed states. The full set of observation-induced mixed states is $\mathcal{R}_{\pi} = \bigcup_{w \in \mathcal{L}} \left\{ \frac{\langle \pi | T^{(w)} \rangle}{\langle \pi | T^{(w)} | 1 \rangle} \right\}$, where $\mathcal{L} \subset \mathcal{A}^*$ is the language of observable words. Although its dynamic is 'meta' in relation to the dynamic of the original HMM, the MSP itself is also a HMM with labeled transition operators $W^{(x)}$ summing to the row-stochastic and generically nondiagonalizable mixed-state-to-state transition dynamic W.

Recall that the *meromorphic functional calculus* prescribes the construction of a new operator from an arbitrary function f of a linear operator A as:

$$f(A) = \sum_{\lambda \in \Lambda_A} \sum_{m=0}^{\nu_{\lambda}-1} A_{\lambda,m} \frac{1}{2\pi i} \oint_{C_{\lambda}} \frac{f(z)}{(z-\lambda)^{m+1}} dz ,$$

where Λ_A is the set of A's eigenvalues, ν_{λ} is the index of the eigenvalue λ (i.e., the size of the largest Jordan block

associated with λ), $z \in \mathbb{C}$, C_{λ} is a positively oriented Jordan curve around λ that includes no other singularities besides possibly at λ itself, and $A_{\lambda,m} = A_{\lambda}(\lambda I - A)^m$ where I is the identity operator and A_{λ} is the spectral projection operator, more simply known as the eigenprojector. The eigenprojector can be constructed as:

$$A_{\lambda} = \sum_{k=1}^{n} \sum_{m=1}^{m_k} \delta_{\lambda,\lambda_k} \left| \lambda_k^{(m)} \right\rangle \left\langle \lambda_k^{(m_k+1-m)} \right| ,$$

where $|\lambda_k^{(m+1)}\rangle$ and $\langle\lambda_k^{(m+1)}|$ are the $m^{\rm th}$ generalized right and left eigenvectors, respectively, of the $k^{\rm th}$ Jordan chain such that:

$$(A - \lambda_k I) |\lambda_k^{(m+1)}\rangle = |\lambda_k^{(m)}\rangle$$

and:

$$\lambda_k^{(m+1)} | (A - \lambda_k I) = \langle \lambda_k^{(m)} | ,$$

for $0 \leq m \leq m_k - 1$, where $|\lambda_j^{(0)}\rangle = \vec{0}$ and $\langle \lambda_j^{(0)}| = \vec{0}$. Specifically, $|\lambda_k^{(1)}\rangle$ and $\langle \lambda_k^{(1)}|$ are conventional right and left eigenvectors, respectively—these are sometimes written more simply as $|\lambda\rangle$ and $\langle \lambda|$, when the algebraic multiplicity $a_{\lambda} = 1$ and there is thus no risk of ambiguity. After imposing normalization, we find that:

$$\langle \lambda_j^{(m)} | \lambda_k^{(n)} \rangle = \delta_{j,k} \delta_{m+n,m_k+1} \; .$$

Perhaps counterintuitively, this implies for example that the most generalized right eigenvectors are dual to the least generalized left eigenvectors.

Usefully, Part I also demonstrated how eigenvalues can be read off from graph motifs and how eigenvectors and spectral projection operators can be built up hierarchically. Part I should be consulted for further details.

B. Outline of main results

With Part I's toolset laid out, Part II now derives the promised closed-form complexities of a process. Section §II investigates the range of possible behaviors for correlation and myopic uncertainty via convergence to asymptotic correlation and asymptotic entropy rates. Section §III then considers measures related to accumulating quantities during the transient relaxation to synchronization. Section §IV introduces closed-form expressions for a wide range of complexity measures in terms of the spectral decomposition of a process' dynamic. It also introduces complexity spectra and highlights common simplifications for special cases, such as almost diagonalizable dynamics. The *excess entropy spectrum* should be an especially useful diagnostic tool since it shows not only the extent but also the timescales over which the past can yield predictions about the future. Additionally, we discover a new timescale—the *symmetry collapse index*—reflective of computational structures in infinite-Markov-order processes. Section §V gives a new kind of signal analysis in terms of coronal spectrograms. A suite of examples in §VI and §VII ground the theoretical developments and are complemented with an in-depth pedagogical example worked out in App. §A. Finally, we conclude with a brief retrospective of Parts I and II and give an eye towards future applications.

II. CORRELATION AND MYOPIC UNCERTAINTY

Using Part I's methods, our first step is to solve for the *autocorrelation function*:

$$\gamma(L) = \left\langle \overline{X}_t X_{t+L} \right\rangle_t \tag{1}$$

and the *myopic uncertainty* or finite-history Shannon entropy rate:

$$h_{\mu}(L) = \mathrm{H}[X_L|X_{1:L}]$$
 (2)

While the first of these indicates how a time-series correlates with itself across time, the second indicates how much new information is generated by a single observation beyond what could be inferred from the last L observations. A comparison is informative. We then determine the asymptotic correlation and myopic uncertainty from the resulting finite-L expressions.

A. Nonasymptotics

A central result in Part I was the spectral decomposition of powers of a linear operator A, even if that operator is nondiagonalizable. Recall that for any $L \in \mathbb{C}$:

$$A^{L} = \left[\sum_{\substack{\lambda \in \Lambda_{A} \\ \lambda \neq 0}} \sum_{m=0}^{\nu_{\lambda}-1} {L \choose m} \lambda^{L-m} A_{\lambda,m} \right] + \left[0 \in \Lambda_{A}\right] \sum_{m=0}^{\nu_{0}-1} \delta_{L,m} A_{0} A^{m} , \qquad (3)$$

where $\binom{L}{m}$ is the generalized binomial coefficient:

$$\binom{L}{m} = \frac{1}{m!} \prod_{n=1}^{m} (L-n+1) , \qquad (4)$$

 $\binom{L}{0} = 1$, and $[0 \in \Lambda_A]$ is the *Iverson bracket*. The latter takes on value 1 if zero is an eigenvalue of A and 0 if not.

In light of this, the autocorrelation function $\gamma(L)$ is simply a superposition of weighted eigen-contributions. Part I showed that Eq. (1) has the operator expression:

$$\gamma(L) = \langle \pi \overline{\mathcal{A}} | T^{|L|-1} | \mathcal{A} \mathbf{1} \rangle ,$$

where T is the transition dynamic of any HMM presentation of the process, \mathcal{A} is the output symbol alphabet, and we defined the row vector:

$$\langle \pi \overline{\mathcal{A}} | = \langle \pi | \left(\sum_{x \in \mathcal{A}} \overline{x} T^{(x)} \right)$$

and the column vector:

$$|\mathcal{A}\mathbf{1}\rangle = \left(\sum_{x\in\mathcal{A}} xT^{(x)}\right)|\mathbf{1}\rangle$$
.

Substituting Part I's spectral decomposition of matrix powers, Eq. (3) above, directly leads to the spectral decomposition of $\gamma(L)$ for nonzero integer L:

$$\gamma(L) = \sum_{\substack{\lambda \in \Lambda_T \\ \lambda \neq 0}} \sum_{m=0}^{\nu_{\lambda}-1} \langle \pi \overline{\mathcal{A}} | T_{\lambda,m} | \mathcal{A} \mathbf{1} \rangle \binom{|L|-1}{m} \lambda^{|L|-1-m} + [0 \in \Lambda_T] \sum_{m=0}^{\nu_0-1} \langle \pi \overline{\mathcal{A}} | T_0 T^m | \mathcal{A} \mathbf{1} \rangle \, \delta_{|L|-1,m} \quad (5)$$
$$= \gamma_{\leadsto}(L) + \gamma_{\multimap}(L) \, . \tag{6}$$

We denote the persistent first term of Eq. (5) as γ_{\rightarrow} , and note that it can be expressed:

$$\begin{split} \gamma_{\leadsto}(L) &= \langle \pi \overline{\mathcal{A}} | \, T^{\mathcal{D}} T T^{|L|-1} \, | \mathcal{A} \mathbf{1} \rangle \\ &= \langle \pi \overline{\mathcal{A}} | \, T^{\mathcal{D}} T^{|L|} \, | \mathcal{A} \mathbf{1} \rangle \ , \end{split}$$

where $T^{\mathcal{D}}$ is T's Drazin inverse. We denote the ephemeral second term as $\gamma_{-\circ}$, which can be written as:

$$\gamma_{-\circ}(L) = \langle \pi \overline{\mathcal{A}} | T_0 T^{|L|-1} | \mathcal{A} \mathbf{1} \rangle$$

where T_0 is the eigenprojector associated with the eigenvalue of zero; $T_0 = \mathbf{0}$ if $0 \notin \Lambda_T$.

From Eq. (5), it is now apparent that the index of T's zero eigenvalue gives a finite-horizon contribution $(\gamma_{-\circ})$ to the autocorrelation function. Beyond index ν_0 of T, the only L-dependence comes via a weighted sum of terms of the form $\binom{|L|-1}{m}$ —polynomials in L times decaying exponentials. The set $\{\langle \pi \overline{\mathcal{A}} | T_{\lambda,m} | \mathcal{A} \mathbf{1} \rangle\}$ simply weights the amplitudes of these contributions. In the familiar diagonalizable case, the behavior of autocorrelation is simply a sum of decaying exponentials $\lambda^{|L|}$.

Similarly, in light of Part I's expression for the myopic

entropy rate in terms of the MSP—starting in the initial unsynchronized mixed-state π and evolving the state of uncertainty via the observation-induced MSP transition dynamic W:

$$h_{\mu}(L) = \langle \delta_{\pi} | W^{L-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle , \qquad (7)$$

where

$$|\mathrm{H}(W^{\mathcal{A}})\rangle \equiv -\sum_{\eta \in \boldsymbol{\mathcal{R}}_{\pi}} |\delta_{\eta}\rangle \sum_{x \in \mathcal{A}} \langle \delta_{\eta} | W^{(x)} | \mathbf{1} \rangle \log_{2} \langle \delta_{\eta} | W^{(x)} | \mathbf{1} \rangle$$

is simply the column vector whose i^{th} entry is the entropy of transitioning from the i^{th} state of S-MSP—and its spectral decomposition of A^L , we find the most general spectral decomposition of the myopic entropy rates $h_{\mu}(L)$ to be:

$$h_{\mu}(L) = \sum_{\substack{\lambda \in \Lambda_{W} \\ \lambda \neq 0}} \sum_{m=0}^{\nu_{\lambda}-1} \langle \delta_{\pi} | W_{\lambda,m} | \mathcal{H}(W^{\mathcal{A}}) \rangle \binom{L-1}{m} \lambda^{L-1-m} + \left[0 \in \Lambda_{W} \right] \sum_{m=0}^{\nu_{0}-1} \delta_{L-1,m} \langle \delta_{\pi} | W_{0}W^{m} | \mathcal{H}(W^{\mathcal{A}}) \rangle$$
(8)
$$= h_{\rightsquigarrow}(L) + h_{\multimap}(L) .$$
(9)

We denote the persistent first term of Eq. (8) as h_{\rightarrow} , and note that it can be expressed directly as:

$$h_{\rightsquigarrow}(L) = \langle \delta_{\pi} | W^{\mathcal{D}} W W^{L-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle$$
$$= \langle \delta_{\pi} | W^{\mathcal{D}} W^{L} | \mathbf{H}(W^{\mathcal{A}}) \rangle ,$$

where $W^{\mathcal{D}}$ is the Drazin inverse of the mixed-stateto-state net transition dynamic W. We denote the ephemeral second term as $h_{-\infty}$, which can be written as:

$$h_{\multimap}(L) = \langle \delta_{\pi} | W_0 W^{L-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle$$

From Eq. (8), we see that the index of W's zero eigenvalue gives a finite horizon contribution $(h_{-\circ})$ to the myopic entropy rate. Beyond index ν_0 of W, the only L-dependence comes via a weighted sum of terms of the form $\binom{L-1}{m}\lambda^{L-1-m}$ —polynomials in L times decaying exponentials. The set $\{\langle \delta_{\pi} | W_{\lambda,m} | \mathcal{H}(W^{\mathcal{A}}) \rangle\}$ weights the amplitudes of these contributions.

For stationary processes we anticipate that, for all $\zeta \in \{\lambda \in \Lambda_W : |\lambda| = 1, \lambda \neq 1\}, \langle \delta_\pi | W_{\zeta} = \mathbf{0}$ and thus $\langle \delta_\pi | W_{\zeta} | \mathbf{H}(W^{\mathcal{A}}) \rangle = 0$. Hence, we can save ourselves from superfluous calculation by excluding the nonunity eigenvalues on the unit circle, when calculating the myopic entropy rate for stationary processes. In the diagonalizable case, again, its behavior is simply a sum of decaying exponentials λ^L .

In practice, $\gamma_{-\infty}$ often vanishes, whereas $h_{-\infty}$ is often

nonzero. This practical difference between $\gamma_{-\circ}$ and $h_{-\circ}$ stems from the difference between typical graph structures of the respective dynamics. For a stationary process' generic transition dynamic, zero eigenvalues (and so $\nu_0(T)$ of T) typically arise from hidden symmetries in the dynamic. In contrast, the MSP of a generic transition dynamic often has tree-like ephemeral structures that are primarily responsible for the zero eigenvalues (and $\nu_0(W)$). Nevertheless, despite their practical typical differences, the same mathematical structures appear and contribute to the most general behavior of each of these cascading quantities.

The breadth of qualitative behaviors shared by autocorrelation and myopic entropy rate is common to the solution of all questions that can be reformulated as a cascading hidden linear dynamic; the myopic state uncertainty $\mathcal{H}^+(L)$ is just one of many other examples. As we have already seen, however, different measures of a process reflect signatures of different linear operators.

Next, we explore similarities in the qualitative behavior of asymptotics and discuss the implications for correlation and entropy rate.

B. Asymptotic correlation

The spectral decomposition reveals that the autocorrelation converges to a constant value as $L \to \infty$, unless T has eigenvalues on the unit circle besides unity itself. This holds if index ν_0 is finite, which it is for all processes generated by finite-state HMMs and also many infinite-state HMMs. If unity is the sole eigenvalue with magnitude one, then all other eigenvalues have magnitude less than unity and their contributions vanish for large enough L. Explicitly, if $\nu_0(T) < \infty$ and $\operatorname{argmax}_{\lambda \in \Lambda_T} |\lambda| = \{1\}$, then:

$$\lim_{L \to \infty} \gamma(L)$$

$$= \lim_{L \to \infty} \sum_{\substack{\lambda \in \Lambda_T \\ \lambda \neq 0}} \sum_{m=0}^{\nu_{\lambda}-1} \langle \pi \overline{\mathcal{A}} | T_{\lambda,m} | \mathcal{A} \mathbf{1} \rangle {\binom{L-1}{m}} \lambda^{L-1-m}$$

$$= \langle \pi \overline{\mathcal{A}} | T_1 | \mathcal{A} \mathbf{1} \rangle$$

$$= \langle \pi \overline{\mathcal{A}} | \mathbf{1} \rangle \langle \pi | \mathcal{A} \mathbf{1} \rangle$$

$$= \left| \sum_{x \in \mathcal{A}} x \Pr(x) \right|^2$$

$$= |\langle x \rangle|^2 .$$

This used the fact that $\nu_1 = 1$ and that $T_1 = |\mathbf{1}\rangle \langle \pi |$ for an ergodic process. This confirms the expectation that $\langle X_t \overline{X}_{t+L} \rangle_t - \langle X_t \rangle_t \langle \overline{X}_t \rangle_t \to 0$ as $L \to \infty$ if $\operatorname{argmax}_{\lambda \in \Lambda_T} |\lambda| = \{1\}$, since the random variables (X_t) If other eigenvalues in Λ_T besides unity lie on the unit circle, then the autocorrelation approaches a periodic sequence as L gets large. This latter case includes not only deterministically periodic processes, but also stochastic processes with periodic randomness.

C. Asymptotic entropy rate

By the Perron–Frobenius theorem, $\nu_{\lambda} = 1$ for all eigenvalues of W on the unit circle. Hence, in the limit of $L \to \infty$, we obtain the asymptotic entropy rate for any stationary process:

$$h_{\mu} \equiv \lim_{L \to \infty} h_{\mu}(L) \tag{10}$$

$$= \lim_{L \to \infty} \sum_{\substack{\lambda \in \Lambda_W \\ |\lambda|=1}} \lambda^{L-1} \left\langle \delta_{\pi} \right| W_{\lambda} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle$$
(11)

$$= \langle \delta_{\pi} | W_1 | \mathcal{H}(W^{\mathcal{A}}) \rangle , \qquad (12)$$

since, for stationary processes, $\langle \delta_{\pi} | W_{\zeta} = \mathbf{0}$ for all $\zeta \in \{\lambda \in \Lambda_W : |\lambda| = 1, \lambda \neq 1\}$. For nonstationary processes, the limit may not exist, but h_{μ} may still be found in a suitable sense as a function of time. If the process has only one stationary distribution over mixed states, then $W_1 = |\mathbf{1}\rangle \langle \pi_W |$ and we have:

$$h_{\mu} = \langle \pi_W | \mathbf{H}(W^{\mathcal{A}}) \rangle \quad , \tag{13}$$

where π_W is the stationary distribution over W's states, found either from $\langle \pi_W | = \langle \delta_\pi | W_1$ or from solving $\langle \pi_W | W = \langle \pi_W |$.

A simple but interesting example of when ergodicity does *not* hold is the multi-armed bandit problem [10, 11]. In this, a realization is drawn from an ensemble of differently biased coins or, for that matter, over any other collection of IID processes. More generally, there can be many distinct memoryful stationary components from which a given realization is sampled, according to some probability distribution. With many attracting components we have the stationary mixed-state eigenprojector $W_1 = \sum_{k=1}^{a_1} |1_k\rangle \langle 1_k|$, with $\langle 1_j |1_k \rangle = \delta_{j,k}$, where the algebraic multiplicity $a_1(T) = a_1(W)$ of the '1' eigenvalue is the number of attracting components. The entropy rate becomes:

$$h_{\mu} = \sum_{k=1}^{a_1} \langle \delta_{\pi} | 1_k \rangle \langle 1_k | \mathbf{H}(W^{\mathcal{A}}) \rangle \tag{14}$$

$$= \left\langle h_{\mu}^{(\text{component }k)} \right\rangle_k \,. \tag{15}$$

Above, $\langle \delta_{\pi} | 1_k \rangle$ is the probability of ending up in compo-

nent k, while $\langle 1_k | \mathcal{H}(W^{\mathcal{A}}) \rangle$ is component k's entropy rate. Thus, if nonergodic, the *process'* entropy rate may not be the same as the entropy of any particular realization. Rather, the process' entropy rate is a weighted average of those for the ensemble of sequences constituting the process.

For unifilar \mathcal{M} , the topology, transition probabilities, and stationary distribution over the recurrent states are the same for both \mathcal{M} and its S-MSP. Hence, for unifilar \mathcal{M} we have:

$$h_{\mu} = \langle \pi_W | \mathbf{H}(W^{\mathcal{A}}) \rangle$$
$$= \langle \pi | H(T^{\mathcal{A}}) \rangle .$$
(16)

One can easily show that Eq. (16) is equivalent to the well-known closed-form expression for h_{μ} for unifilar presentations:

$$\langle \pi | H(T^{\mathcal{A}}) \rangle = -\sum_{\sigma \in \mathcal{S}} \Pr(\sigma) \sum_{\substack{x \in \mathcal{A} \\ \sigma' \in \mathcal{S}}} T^{(x)}_{\sigma,\sigma'} \log_2(T^{(x)}_{\sigma,\sigma'}) . \quad (17)$$

For nonunifilar presentations, however, we must use the more general result of Eq. (13). This is similar to the calculation in Eq. (17), but must be performed over the recurrent states of a mixed-state presentation, which may be countable or uncountable.

III. ACCUMULATED TRANSIENTS FOR DIAGONALIZABLE DYNAMICS

In the diagonalizable case, autocorrelation, myopic entropy rate, and myopic state uncertainty reduce to a sum of decaying exponentials. Correspondingly, we can find the power spectrum, excess entropy, and synchronization information respectively via geometric progressions.

For example, if W is diagonalizable and has no zero eigenvalue, then the myopic entropy rate reduces to:

$$h_{\mu}(L) = \sum_{\lambda \in \Lambda_{W}} \langle \delta_{\pi} | W_{\lambda} | \mathbf{H}(W^{\mathcal{A}}) \rangle \lambda^{L-1}$$

= $\langle \delta_{\pi} | W_{1} | \mathbf{H}(W^{\mathcal{A}}) \rangle + \sum_{\substack{\lambda \in \Lambda_{W} \\ |\lambda| < 1}} \lambda^{L-1} \langle \delta_{\pi} | W_{\lambda} | \mathbf{H}(W^{\mathcal{A}}) \rangle,$

where $\langle \delta_{\pi} | W_1 | \mathcal{H}(W^{\mathcal{A}}) \rangle$ is identifiable as the entropy rate h_{μ} .

It then follows that the *excess entropy*, which is the mutual information $\mathbf{E} = \mathbf{I}[\overleftarrow{X}; \overrightarrow{X}]$ between the past and the future—and thus how much of the future can be pre-

7

dicted from the past—is:

$$\mathbf{E} \equiv \sum_{L=1}^{\infty} [h_{\mu}(L) - h_{\mu}]$$

$$= \sum_{L=1}^{\infty} \sum_{\substack{\lambda \in \Lambda_{W} \\ |\lambda| < 1}} \lambda^{L-1} \langle \delta_{\pi} | W_{\lambda} | \mathbf{H}(W^{\mathcal{A}}) \rangle$$

$$= \sum_{\substack{\lambda \in \Lambda_{W} \\ |\lambda| < 1}} \langle \delta_{\pi} | W_{\lambda} | \mathbf{H}(W^{\mathcal{A}}) \rangle \sum_{\substack{L=1 \\ = \sum_{L=0}^{\infty} \lambda^{L} = \frac{1}{1-\lambda}}}^{\infty} \lambda^{L-1}$$

$$= \sum_{\substack{\lambda \in \Lambda_{W} \\ |\lambda| < 1}} \frac{1}{1-\lambda} \langle \delta_{\pi} | W_{\lambda} | \mathbf{H}(W^{\mathcal{A}}) \rangle .$$
(18)

Note that larger eigenvalues (closer to unity magnitude) drive the denominator $1 - \lambda$ closer to zero and, thus, increase $\frac{1}{1-\lambda}$. Hence, larger eigenvalues—controlling modes of the mixed-state transition matrix that decay slowly—have the potential to contribute most to excess entropy. Small eigenvalues—quickly decaying modes—do not contribute significantly. Putting aside the language of eigenvalues, one can paraphrase: slowly decaying transient behavior (of the distribution of distributions over process states) has the most potential to make a process appear complex.

Continuing, the *transient information*, used in the context of synchronization and distinguishing periodic structures [12], is:

$$\begin{aligned} \mathbf{T} &\equiv \sum_{L=1}^{\infty} L \left[h_{\mu}(L) - h_{\mu} \right] \\ &= \sum_{L=1}^{\infty} \sum_{\substack{\lambda \in \Lambda_{W} \\ |\lambda| < 1}} L \lambda^{L-1} \left\langle \delta_{\pi} \right| W_{\lambda} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \\ &= \sum_{\substack{\lambda \in \Lambda_{W} \\ |\lambda| < 1}} \left\langle \delta_{\pi} \right| W_{\lambda} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \sum_{\substack{L=1 \\ \ell = 1 \\ \ell = 1$$

We now see that the transient information is very closely related to the excess entropy, differing only via the square in the denominators. This comparison between **E** and **T** closed-form expressions suggests an entire hierarchy of informational quantities based on eigenvalue weighting.

Performing a similar procedure for the synchronization

information \mathbf{S}' [13] shows that:

$$\begin{split} \mathbf{S}' &\equiv \sum_{L=0}^{\infty} \left[\mathcal{H}(L) - \mathcal{H} \right] \\ &= \sum_{L=0}^{\infty} \sum_{\substack{\lambda \in \Lambda_W \\ |\lambda| < 1}} \left\langle \delta_{\pi} \right| W_{\lambda} \left| \mathbf{H}[\eta] \right\rangle \lambda^{L} \\ &= \sum_{\substack{\lambda \in \Lambda_W \\ |\lambda| < 1}} \left\langle \delta_{\pi} \right| W_{\lambda} \left| \mathbf{H}[\eta] \right\rangle \sum_{L=0}^{\infty} \lambda^{L} \\ &= \sum_{\substack{\lambda \in \Lambda_W \\ |\lambda| < 1}} \frac{1}{1 - \lambda} \left\langle \delta_{\pi} \right| W_{\lambda} \left| \mathbf{H}[\eta] \right\rangle , \end{split}$$

where $|\mathbf{H}[\eta] \rangle \equiv \sum_{\eta \in \mathbf{\mathcal{R}}_{\pi}} |\delta_{\eta}\rangle H[\eta]$ is the column vector of entropies associated with each mixed-state.

The expressions reveal a remarkably close relationship between \mathbf{S}' and \mathbf{E} . Define $\langle \cdot | \equiv \sum_{L=0}^{\infty} \langle \delta_{\pi} | W^L$. Then:

$$\left\langle \cdot \right| = \sum_{\substack{\lambda \in \Lambda_W \ |\lambda| < 1}} \frac{1}{1 - \lambda} \left\langle \delta_{\pi} \right| W_{\lambda}$$

The relationship is now made plain:

$$\mathbf{E} = \langle \cdot | \mathbf{H}(W^{\mathcal{A}}) \rangle \text{ and}$$
$$\mathbf{S}' = \langle \cdot | \mathbf{H}[\eta] \rangle .$$

Although a bit more cumbersome, perhaps better intuition emerges if we rewrite $\langle \cdot |$ as $\langle \int \Pr(\eta, L) dL |$.

Again, large eigenvalues—slowly decaying modes of the mixed-state transition matrix—can make the largest contribution to synchronization information; small eigenvalues correspond to quickly decaying modes that do not have the opportunity to contribute. In fact, the potential of large eigenvalues to make large contributions is a recurring theme for many questions one has about a process. Simply stated, long-term behavior—what we often interpret as "complex" behavior—is dominated by a process's largest-eigenvalue modes.

That said, a word of warning is in order. Although large-eigenvalue modes have the most *potential* to make contributions to a process's complexity, the actual set of largest contributors also depends strongly on the amplitudes $\{\langle \delta_{\pi} | W_{\lambda} | \dots \rangle\}$, where $|\dots\rangle$ is some quantifier vector of interest; e.g., $|\dots\rangle = |\mathrm{H}[\eta]\rangle$, $|\dots\rangle = |\mathrm{H}(W^{\mathcal{A}})\rangle$, or $|\dots\rangle = |\mathbf{1}\rangle$.

Hence, there is as-yet unanticipated similarity between \mathbf{E} and \mathbf{T} and another between \mathbf{E} and \mathbf{S}' —at least assuming diagonalizability. We would like to know the relationships between these quantities more generally. However, deriving the general closed-form expressions for accumulated transients is not tractable via the current approach.

Rather, to derive the general results, we deploy the meromorphic functional calculus directly at an elevated level, as we now demonstrate.

IV. EXACT COMPLEXITIES AND COMPLEXITY SPECTRA

We now derive the most general closed-form solutions for several complexity measures, from which expressions for related measures follow straightforwardly. This includes an expression for the past-future mutual information or excess entropy, identifying two distinct persistent and transient components, and a novel extension of excess entropy to temporal frequency spectra components. We also give expressions for the synchronization information and power spectra. We explicitly address the class a common one we argue—of almost diagonalizable dynamics. The section finishes by highlighting finite-order Markov order processes that, rather than being simpler than infinite Markov order processes, introduce technical complications that must be addressed.

Before carrying this out, we define several useful objects. Let $\rho(A)$ be the spectral radius of matrix A:

$$\rho(A) = \max_{\lambda \in \Lambda_A} |\lambda| \; .$$

For stochastic W, since $\rho(W) = 1$, let $\Lambda_{\rho(W)}$ denote the set of eigenvalues with unity magnitude:

$$\Lambda_{\rho(W)} = \{\lambda \in \Lambda_W : |\lambda| = 1\} .$$

We also define:

$$Q \equiv W - W_1 \tag{19}$$

and

$$Q \equiv W - \sum_{\lambda \in \Lambda_{\rho(W)}} \lambda W_{\lambda} .$$
 (20)

Eigenvalues with unity magnitude that are not themselves unity correspond to perfectly periodic cycles of the state-transition dynamic. By their very nature, such cycles are restricted to the recurrent states. Moreover, we expect the projection operators associated with these cycles to have no net overlap with the start-state of the MSP. So, we expect:

$$\langle \delta_{\pi} | W_{\lambda} = \vec{0} , \qquad (21)$$

for all $\lambda \in \Lambda_{\rho(W)} \setminus \{1\}$. Hence:

$$\langle \delta_{\pi} | Q^L = \langle \delta_{\pi} | Q^L .$$
 (22)

We will also use the fact that, since $\rho(Q) < 1$:

$$\sum_{L=0}^{\infty} \mathcal{Q}^L = (I - \mathcal{Q})^{-1} ;$$

and furthermore:

$$\langle \delta_{\pi} | (I - \mathcal{Q})^{-1} = \langle \delta_{\pi} | (I - Q)^{-1} ,$$

as a consequence of Eq. (21) and our spectral decomposition.

Having seen complexity measures associated with prediction all take on a similar form in terms of the S-MSP state-transition matrix, we expect to encounter similar forms for generically nondiagonalizable state-transition dynamics.

A. Excess entropy

We are now ready to develop the excess entropy in full generality. Our tools turn this into a direct calculation. We find:

$$\begin{split} \mathbf{E} &\equiv \sum_{L=1}^{\infty} \left[h_{\mu}(L) - h_{\mu} \right] \\ &= \sum_{L=1}^{\infty} \left[\langle \delta_{\pi} | W^{L-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle - \langle \delta_{\pi} | W_{1} | \mathbf{H}(W^{\mathcal{A}}) \rangle \right] \\ &= \sum_{L=0}^{\infty} \left[\langle \delta_{\pi} | W^{L} | \mathbf{H}(W^{\mathcal{A}}) \rangle - \langle \delta_{\pi} | W_{1} | \mathbf{H}(W^{\mathcal{A}}) \rangle \right] \\ &= \sum_{L=0}^{\infty} \langle \delta_{\pi} | \left[(\underbrace{W - W_{1}}_{\equiv Q})^{L} - \delta_{L,0} W_{1} \right] | \mathbf{H}(W^{\mathcal{A}}) \rangle \\ &= - \underbrace{\langle \delta_{\pi} | W_{1} | \mathbf{H}(W^{\mathcal{A}}) \rangle}_{=h_{\mu}} + \sum_{L=0}^{\infty} \underbrace{\langle \delta_{\pi} | Q^{L}}_{=\langle \delta_{\pi} | Q^{L}} | \mathbf{H}(W^{\mathcal{A}}) \rangle \\ &= \langle \delta_{\pi} | \left(\sum_{L=0}^{\infty} Q^{L} \right) | \mathbf{H}(W^{\mathcal{A}}) \rangle - h_{\mu} \\ &= \langle \delta_{\pi} | \left(I - Q \right)^{-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle - h_{\mu} \\ &= \langle \delta_{\pi} | \left(I - Q \right)^{-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle - h_{\mu} \end{split}$$

Note that $(I - Q)^{-1} = \text{inv}(I - Q)$ here, since unity is not an eigenvalue of Q. Indeed, the unity eigenvalue was explicitly extracted from the former matrix to make an invertible expression.

For an ergodic process, where $W_1 = |\mathbf{1}\rangle \langle \pi_W |$, this becomes:

$$\mathbf{E} = \langle \delta_{\pi} | \left(I - W + |\mathbf{1}\rangle \langle \pi_{W} | \right)^{-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle - h_{\mu} . \quad (23)$$

Computationally, Eq. (23) is wonderfully useful. How-

ever, the subtraction of h_{μ} is at first mysterious. Especially so, when compared to the compact result for the excess-entropy spectral decomposition in the diagonalizable case given by Eq. (18).

Let's explore this. Recall that Ref. [2] showed:

$$(I-T)^{\mathcal{D}} = [I - (T - T_1)]^{-1} - T_1 ,$$
 (24)

for any stochastic matrix T, where T_1 is the projection operator associated with eigenvalue $\lambda = 1$. From this, we see that the general solution for **E** takes on its most elegant form in terms of the Drazin inverse of I - W:

$$\mathbf{E} = \langle \delta_{\pi} | (I - Q)^{-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle - h_{\mu}$$

= $\langle \delta_{\pi} | [(I - Q)^{-1} - W_1] | \mathbf{H}(W^{\mathcal{A}}) \rangle$
= $\langle \delta_{\pi} | (I - W)^{\mathcal{D}} | \mathbf{H}(W^{\mathcal{A}}) \rangle$. (25)

Recall too Part I's explicit spectral decomposition:

$$(I-T)^{\mathcal{D}} = \sum_{\lambda \in \Lambda_T \setminus \{1\}} \sum_{m=0}^{\nu_{\lambda}-1} \frac{1}{(1-\lambda)^{m+1}} T_{\lambda,m} , \qquad (26)$$

which uses the companion operators $T_{\lambda,m}$ from there. From this and Eq. (25), we see that the past-future mutual information—the amount of the future that is *predictable* from the past—has the general spectral decomposition:

$$\mathbf{E} = \sum_{\lambda \in \Lambda_W \setminus \{1\}} \sum_{m=0}^{\nu_{\lambda}-1} \frac{1}{(1-\lambda)^{m+1}} \left\langle \delta_{\pi} \right| W_{\lambda,m} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle.$$
(27)

B. Persistent excess

In light of Eq. (9), we see that there are two qualitatively distinct contributions to the excess entropy $\mathbf{E} = \mathbf{E}_{\rightarrow \rightarrow} + \mathbf{E}_{\rightarrow \circ}$. One comprises the persistent leaky contributions from all L:

$$\mathbf{E}_{\leadsto} \equiv \sum_{L=1}^{\infty} \left[h_{\leadsto}(L) - h_{\mu} \right]$$
$$= \left\langle \delta_{\pi} \right| W^{\mathcal{D}} W (I - W)^{\mathcal{D}} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle$$

and the other is a completely ephemeral piece that contributes only up to W's zero-eigenvalue index ν_0 :

$$\begin{split} \mathbf{E}_{-\circ} &\equiv \sum_{L=1}^{\infty} h_{-\circ}(L) \\ &= \sum_{L=1}^{\nu_0} h_{-\circ}(L) \\ &= \left\langle \delta_{\pi} \right| W_0 (I - W)^{\mathcal{D}} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \;. \end{split}$$

C. Excess entropy spectrum

Equation (25) immediately suggests that we generalize the excess entropy, a scalar complexity measure, to a complexity *function* with continuous part defined in terms of the resolvent—say, via introducing the complex variable z:

$$E(z) = \langle \delta_{\pi} | (zI - W)^{-1} | \mathbf{H}(W^{\mathcal{A}}) \rangle$$

Such a function not only monitors how much of the future is predictable, but also reveals the time scales of interdependence between the predictable features within the observations. Directly taking the z-transform of $h_{\mu}(L)$ comes to mind, but this requires tracking both real and imaginary parts or, alternatively, both magnitude and complex phase. To ameliorate this, we employ a transform of a closely related function that contains the same information.

Before doing so, we should briefly note that ambiguity surrounds the appropriate excess-entropy generalization. There are many alternate measures that approach the excess entropy as frequency goes to zero. For example, directly calculating from the meromorphic functional calculus, letting $z = e^{i\omega}$ we find:

$$\lim_{\omega \to 0} \operatorname{Re} \left\langle \delta_{\pi} \right| \left(e^{i\omega} I - W \right)^{-1} \left| \operatorname{H}(W^{\mathcal{A}}) \right\rangle = \mathbf{E} - \frac{1}{2} h_{\mu} \ .$$

We are challenged, however, to interpret the fact that $\operatorname{Re} \langle \delta_{\pi} | (e^{i\omega}I - W)^{-1} | \mathrm{H}(W^{\mathcal{A}}) \rangle + \frac{1}{2}h_{\mu}$ is not necessarily positive at all frequencies. Another direct calculation shows that:

$$\lim_{\omega \to 0} \operatorname{Re} \left\langle \delta_{\pi} \right| e^{i\omega} (e^{i\omega} I - W)^{-1} \left| \operatorname{H}(W^{\mathcal{A}}) \right\rangle = \mathbf{E} + \frac{1}{2} h_{\mu} .$$

Enticingly, Re $\langle \delta_{\pi} | e^{i\omega} (e^{i\omega}I - W)^{-1} | \mathcal{H}(W^{\mathcal{A}}) \rangle - \frac{1}{2}h_{\mu}$ appears to be positive over all frequencies for all examples checked. It is not immediately clear which, if either, is the appropriate generalization, though. Fortunately, the Fourier transform of a two-sided myopic-entropy convergence function makes our upcoming definition of $\mathcal{E}(\omega)$ interpretable and of interest in its own right.

Let dh be the *two-sided myopic entropy convergence* function defined by:

$$\mathrm{dh}(L) = \begin{cases} \mathrm{H}[X_0|X_{-|L|+1:0}] & \text{for } L < 0 \ , \\ \log_2 |\mathcal{A}| & \text{for } L = 0 \ , \text{and} \\ \mathrm{H}[X_0|X_{1:L}] & \text{for } L > 0 \ . \end{cases}$$

For stationary processes, it is easy to show that $H[X_0|X_{-L+1:0}] = H[X_0|X_{1:L}]$, with the result that dh

is a symmetric function. Moreover, dh then simplifies to:

$$\mathrm{dh}(L) = h_{\mu}(|L|) ,$$

where $h_{\mu}(0) \equiv \log_2 |\mathcal{A}|$ and, as before, $h_{\mu}(L) = H[X_L|X_{1:L}]$ for $L \geq 1$ with $h_{\mu}(1) = H[X_1]$.

The symmetry of the two-sided myopic entropy convergence function dh guarantees that its Fourier transform is also real and symmetric. Explicitly, the continuous part of the Fourier transform turns out to be:

$$\widetilde{\mathrm{dh}}_{\mathrm{c}}(\omega) = \mathbf{R} + 2\mathrm{Re}\left\langle \delta_{\pi} \right| \left(e^{i\omega}I - W \right)^{-1} |\mathrm{H}(W^{\mathcal{A}})\rangle$$

a strictly real and symmetric function of the angular frequency ω . Here, **R** is the redundancy of the alphabet $\mathbf{R} \equiv \log_2 |\mathcal{A}| - h_{\mu}$, as in Ref. [12].

The transform dh also has a discrete impulsive component. For stationary processes this consists solely of the Dirac δ -function at zero frequency:

$$\widetilde{\mathrm{dh}}_{\mathrm{d}}(\omega) = 2\pi h_{\mu} \sum_{k \in \mathbb{Z}} \delta(\omega + 2\pi k) \; .$$

Recall that the Fourier transform of a discrete-domain function is 2π -periodic in the angular frequency ω . This δ -function is associated with the nonzero offset of the entropy convergence curve of positive-entropy-rate processes. The full transform is:

$$\widetilde{\mathrm{dh}}(\omega) = \widetilde{\mathrm{dh}}_{\mathrm{c}}(\omega) + \widetilde{\mathrm{dh}}_{\mathrm{d}}(\omega) \,.$$

Direct calculation using Ref. [2]'s meromorphic functional calculus shows that:

$$\lim_{\omega \to 0} \operatorname{Re} \left\langle \delta_{\pi} \right| \left(e^{i\omega} I - W \right)^{-1} \left| \operatorname{H}(W^{\mathcal{A}}) \right\rangle = \mathbf{E} - \frac{1}{2} h_{\mu} \,. \quad (28)$$

This motivates introducing the excess-entropy spectrum $\mathcal{E}(\omega)$:

$$\mathcal{E}(\omega) \equiv \frac{1}{2} \left(\widehat{\mathrm{dh}}(\omega) - \mathbf{R} + h_{\mu} \right)$$
(29)
$$= \operatorname{Re} \left\langle \delta_{\pi} \right| \left(e^{i\omega} I - W \right)^{-1} \left| \mathrm{H}(W^{\mathcal{A}}) \right\rangle + \frac{1}{2} h_{\mu}$$
$$+ \pi h_{\mu} \sum_{k \in \mathbb{Z}} \delta(\omega + 2\pi k) .$$
(30)

The excess-entropy spectrum rather directly displays important frequencies of apparent entropy reduction. For example, leaky period-5 processes have a period-5 signature in the excess entropy spectrum.

As with its predecessors, the excess-entropy spectrum also has a natural decomposition into two qualitatively distinct components:

$$\mathcal{E}(\omega) = \mathcal{E}_{\leadsto}(\omega) + \mathcal{E}_{\multimap}(\omega) \; .$$

The excess-entropy spectrum gives an intuitive and concise summary of the complexities associated with a process' predictability. For example, given a graph of the excess entropy spectrum, the past-future mutual information can be read off as the height of the continuous part of the function as it approaches zero frequency:

$$\mathbf{E} = \lim_{\omega \to 0} \mathcal{E}(\omega)$$
$$= \mathcal{E}_{\mathbf{c}}(\omega = 0) \ .$$

Indeed, the *limit* of zero frequency is necessary due to the δ -function in the Fourier transform at exactly zero frequency:

$$h_{\mu} = \lim_{\epsilon \to 0} \frac{1}{\pi} \int_{-\epsilon}^{\epsilon} \mathcal{E}(\omega) d\omega$$
.

Reflecting on this, the δ -function indicates one of the reasons the excess entropy has been difficult to compute in the past. This also sheds light on the role of the Drazin inverse: It removes the infinite asymptotic accumulation, revealing the transient structure of entropy convergence.

We also have a spectral decomposition of the excessentropy spectrum:

$$\begin{aligned} \mathcal{E}(\omega)_{c} &= \sum_{\lambda \in \Lambda_{W}} \sum_{m=0}^{\nu_{\lambda}-1} \operatorname{Re}\left(\frac{\langle \delta_{\pi} | W_{\lambda,m} | \mathrm{H}(W^{\mathcal{A}}) \rangle}{(e^{i\omega} - \lambda)^{m+1}}\right) \\ &= \sum_{m=0}^{\nu_{0}-1} \cos\left((m+1)\omega\right) \langle \delta_{\pi} | W_{0}W^{m} | \mathrm{H}(W^{\mathcal{A}}) \rangle \\ &+ \sum_{\lambda \in \Lambda_{W}} \sum_{m=0}^{\nu_{\lambda}-1} \operatorname{Re}\left(\frac{\langle \delta_{\pi} | W_{\lambda,m} | \mathrm{H}(W^{\mathcal{A}}) \rangle}{(e^{i\omega} - \lambda)^{m+1}}\right), \end{aligned}$$

where, in the last equality, we assume that W_0 is real. This shows that, in addition to the contribution of typical leaky modes of decay in entropy convergence, the zero-eigenvalue modes contribute uniquely to the excess entropy spectrum. In addition to Lorentzian-like spectral curves contributed by leaky periodicities in the MSP, the excess-entropy spectrum also contains sums of cosines up to a frequency controlled by index ν_0 , which corresponds to the depth of the MSP's nondiagonalizability. This is simply the duration of ephemeral synchronization in the time domain.

D. Synchronization information

Once expressed in terms of the S-MSP transition dynamic, the derivation of the excess synchronization information \mathbf{S}' closely parallels that of the excess entropy, only with a different ket $|\cdot\rangle$ appended. We calculate, as before, finding:

$$\begin{split} \mathbf{S}' &\equiv \sum_{L=0}^{\infty} \left[\mathcal{H}(L) - \mathcal{H} \right] \\ &= \sum_{L=0}^{\infty} \left[\left\langle \delta_{\pi} \right| W^{L} \left| \mathbf{H}[\eta] \right\rangle - \left\langle \delta_{\pi} \right| W_{1} \left| \mathbf{H}[\eta] \right\rangle \right] \\ &= \sum_{L=0}^{\infty} \left\langle \delta_{\pi} \right| \left[\left(\underbrace{W - W_{1}}_{\equiv Q} \right)^{L} - \delta_{L,0} W_{1} \right] \left| \mathbf{H}[\eta] \right\rangle \\ &= - \underbrace{\left\langle \delta_{\pi} \right| W_{1} \left| \mathbf{H}[\eta] \right\rangle}_{=\mathcal{H}} + \sum_{L=0}^{\infty} \underbrace{\left\langle \delta_{\pi} \right| Q^{L}}_{=\left\langle \delta_{\pi} \right| Q^{L}} \left| \mathbf{H}[\eta] \right\rangle \\ &= \left\langle \delta_{\pi} \right| \left(\sum_{L=0}^{\infty} \mathcal{Q}^{L} \right) \left| \mathbf{H}[\eta] \right\rangle - \mathcal{H} \\ &= \left\langle \delta_{\pi} \right| \left(I - \mathcal{Q} \right)^{-1} \left| \mathbf{H}[\eta] \right\rangle - \mathcal{H} \\ &= \left\langle \delta_{\pi} \right| \left(I - Q \right)^{-1} \left| \mathbf{H}[\eta] \right\rangle - \mathcal{H} . \end{split}$$

For an ergodic process where $W_1 = |\mathbf{1}\rangle \langle \pi_W |$, this becomes:

$$\mathbf{S}' = \langle \delta_{\pi} | \left(I - W + |\mathbf{1}\rangle \langle \pi_W | \right)^{-1} |\mathbf{H}[\eta] \rangle - \mathcal{H} .$$
 (31)

From Eq. (24), we see that the general solution for \mathbf{S}' takes on its most elegant form in terms of the Drazin inverse of I - W:

$$\mathbf{S}' = \langle \delta_{\pi} | (I - Q)^{-1} | \mathbf{H}[\eta] \rangle - \mathcal{H}$$

= $\langle \delta_{\pi} | [(I - Q)^{-1} - W_1] | \mathbf{H}[\eta] \rangle$
= $\langle \delta_{\pi} | (I - W)^{\mathcal{D}} | \mathbf{H}[\eta] \rangle$. (32)

From Eq. (32) and Eq. (26), we also see that the excess synchronization information has the general spectral decomposition:

$$\mathbf{S}' = \sum_{\lambda \in \Lambda_W \setminus \{1\}} \sum_{m=0}^{\nu_{\lambda}-1} \frac{1}{(1-\lambda)^{m+1}} \left\langle \delta_{\pi} \right| W_{\lambda,m} \left| \mathbf{H}[\eta] \right\rangle .$$
(33)

Again the form of Eq. (32) suggests generalizing synchronization information from a complexity measure to a complexity function $S(\omega)$. In this case, the result is simply related to the Fourier transform of the two-sided myopic state-uncertainty $\mathcal{H}(L)$.

E. Power spectra

The extended complexity functions, $\mathcal{E}(\omega)$ and $\mathcal{S}(\omega)$ just introduced, give the same intuitive understanding for entropy reduction and synchronization respectively as the power spectrum $P(\omega)$ gives for pairwise correlation. Recall from Part I that the power spectrum can be written as:

j

$$P_{\rm c}(\omega) = \left\langle \left| x \right|^2 \right\rangle + 2 \operatorname{Re} \left\langle \pi \overline{\mathcal{A}} \right| \left(e^{i\omega} I - T \right)^{-1} \left| \mathcal{A} \mathbf{1} \right\rangle \ .$$

We see that $(e^{i\omega}I - T)^{-1}$ is the resolvent of T evaluated along the unit circle $z = e^{i\omega}$ for $\omega \in [0, 2\pi)$. Hence, by Part I's decomposition of the resolvent, the general spectral decomposition of the continuous part of the power spectrum is:

$$P_{\rm c}(\omega) = \left\langle \left| x \right|^2 \right\rangle + 2 \sum_{\lambda \in \Lambda_T} \sum_{m=0}^{\nu_{\lambda}-1} \operatorname{Re} \frac{\left\langle \pi \overline{\mathcal{A}} \right| T_{\lambda,m} \left| \mathcal{A} \mathbf{1} \right\rangle}{(e^{i\omega} - \lambda)^{m+1}} \,.$$

As with $\mathcal{E}(\omega)$ and $\mathcal{S}(\omega)$, all continuous frequency dependence of the power spectrum again lies simply and entirely in the denominator of the above expression.

Analogous to Ref. [14]'s results, the power-spectrum δ functions arise from the eigenvalues of T that lie on the unit circle:

$$P_{\rm d}(\omega) = \sum_{k=-\infty}^{\infty} \sum_{\substack{\lambda \in \Lambda_T \\ |\lambda|=1}} 2\pi \,\delta(\omega - \omega_{\lambda} + 2\pi k) \\ \times \operatorname{Re}\left(\lambda^{-1} \langle \pi \overline{\mathcal{A}} | T_{\lambda} | \mathcal{A}\mathbf{1} \rangle\right) \,,$$

where ω_{λ} is related to λ by $\lambda = e^{i\omega_{\lambda}}$. An extension of the Perron–Frobenius theorem guarantees that the eigenvalues of T on the unit circle have index $\nu_{\lambda} = 1$.

Together, these equations yield structural constraints via particular functional forms that are key to solving the inverse problem of inferring process models from measured data.

F. Almost diagonalizable dynamics

The nondiagonalizability that appears most commonly in prediction metadynamics is of a special form that we call *almost diagonalizable*: when all eigenspaces except one—usually that associated with $\lambda = 0$ —are diagonalizable subspaces. In the current setting, we say that a matrix is *almost diagonalizable* if all of its eigenvalues with magnitude greater than zero have geometric multiplicity equal to their algebraic multiplicity.

Definition 1. W is almost diagonalizable if and only if $g_{\lambda} = a_{\lambda}$ for all $\lambda \in \Lambda_W^{\setminus 0} \equiv \Lambda_W \setminus \{0\}.$

Fortunately, we treat such nondiagonalizability straightforwardly using W^L 's spectral decomposition for singular matrices. First off, Eq. (3) simplifies to:

$$W^L = \sum_{\lambda \in \Lambda_W} \lambda^L W_\lambda + \sum_{m=1}^{\nu_0 - 1} \delta_{L,m} W_0 W^m .$$
 (34)

12

Then, to obtain the projection operators associated with each eigenvalue in $\Lambda_W^{\setminus 0}$ for an almost diagonalizable matrix W, we use Part I's expression for operators with index-one eigenvalues with $\nu_{\lambda} = 1$ for all $\lambda \in \Lambda_W^{\setminus 0}$. Finding:

$$W_{\lambda} = \left(\frac{W}{\lambda}\right)^{\nu_{0}} \prod_{\substack{\zeta \in \Lambda_{V}^{\setminus 0} \\ \zeta \neq \lambda}} \frac{W - \zeta I}{\lambda - \zeta} , \qquad (35)$$

for each $\lambda \in \Lambda_W^{\setminus 0}$. Or, when more convenient in a calculation, we let $\nu_0 \to a_0 - g_0 + 1$ or even $\nu_0 \to a_0$ in Eq. (35), since multiplying W_{λ} by W/λ has no effect.

With the set of projection operators W_{λ} for all $\lambda \in \Lambda_W^{\setminus 0}$ in hand, we can use the fact from Part I that projection operators sum to the identity to determine the projection operator associated with the zero eigenvalue:

$$W_0 = I - \sum_{\lambda \in \Lambda_W^{\setminus 0}} W_\lambda$$
.

This is sometimes simpler and easier to automate than evaluating W_0 via the methods of symbolic inversion and residues or via finding all left and right eigenvectors and generalized eigenvectors.

Almost diagonalizable metadynamics play a prominent role in prediction for both processes of finite Markov order and for the much more general class of processes with broken partial symmetries that can be detected within a finite observation window—the processes of finite symmetry-collapse discussed next.

G. Markov order versus symmetry collapse

What if zero is the *only* eigenvalue in the transient structure of a process' MSP? That is, what if there are no loops in the S-MSP transient structure? The associated processes turn out to have finite Markov order.

For processes with finite Markov order R—such as, those whose support is a *subshift of finite type* [15]—the entropy-rate approximates not only converge but also become equal to the true entropy rate when conditioning on long enough histories. Explicitly, for $\ell \geq R + 1$ [12]:

$$h_{\mu}(\ell) - h_{\mu} = 0 , \qquad (36)$$

or, equivalently, for $L \ge R$:

$$\langle \delta_{\pi} | W^L | \mathbf{H}(W^{\mathcal{A}}) \rangle - \langle \pi_W | \mathbf{H}(W^{\mathcal{A}}) \rangle = 0$$
.

For a finite-order Markov process, all MSP transient states must have identically zero probability after R time steps. The only way to achieve this is if the S-MSP's transient structure is an acyclic directed graph with all probability density flowing away from the unique startstate down to the recurrent component. This means that all eigenvalues associated with the transient states are zero. Moreover, the index of the zero-eigenvalue of the ϵ -machine's S-MSP is equal to the Markov order for finite Markov-order processes. That is, if $\Lambda_W \setminus \Lambda_T = \{0\}$, then:

$$\nu_0(W) = R \; .$$

In contrast, for stochastic processes whose support is a strictly sofic subshift [15], the Markov order diverges, but ν_0 can vanish or be finite or infinite. Yet, in either the finite-type or sofic case, ν_0 still tracks the duration of exact state-space collapse within the transient dynamics of synchronization. This suggests that ν_0 captures the index of broken symmetries for strictly sofic processes, in analogy to the Markov order for subshifts of finite type. The name symmetry-collapse index captures the essence of ν_0 's role in both cases.

Let's explain. In the first ν_0 time steps, symmetries are broken that synchronize an observer to the process. For the simple period-two process ...01010101010... the "symmetry" that is broken is the degeneracy of possible phases—the 0 phase or the 1 phase of the period-2 oscillation. Initially, without making a measurement the two phases are indistinguishable. After a single observation, though, the observer learns the phase and is completely synchronized to the process. Hence, $\nu_0 = 1$ for this order-1 Markov process. Simple periodic processes with larger periods have a longer time before the phase information is fully known; hence, their larger Markov order.

For the more complex strictly sofic processes, there may also be symmetries, such as phase information, that are completely broken within a finite amount of time. However, this is only part of the overall transient metadynamics of synchronization. And so, the symmetries completely broken within the symmetry-collapse epoch occur in addition to lingering state uncertainties about a strictly sofic process. As a practical matter, a process' predictability is often substantially enhanced through the finite epoch of symmetry-collapse. This becomes apparent in the examples to follow.

V. SPECTRAL ANALYSIS VIA CORONAL SPECTROGRAMS

Coronal spectrograms are a broadly useful tool in visualizing complexity spectra, from power spectra to excess entropy spectra. They were recently introduced by Ref. [14] to demonstrate how diffraction patterns of chaotic crystals emanate from the eigenvalue spectrum



(a) How spectra emanate from eigenvalues: Coronal spectrogram (far right) combines a discrete-time process' eigenvalues Λ_T (far left) of the hidden linear dynamic T together with a frequency-dependent function $P(\omega)$ (middle) by wrapping the latter around the unit circle.



(b) Coronated horizon (far right) combines the frequency-dependent function $f(\omega)$ (far left) together with a continuous-time process eigenvalues Λ_A of the hidden linear dynamic by unwrapping the unit circle.

FIG. 1. Spectra and eigenvalues: (a) Coronal spectrogram and (b) coronated horizon.

of the hidden spatial dynamic of stacked modular layers.

Coronal spectrograms display any frequencydependent measure $f(\omega)$ of a process wrapped around the unit circle while showing the eigenvalues Λ_T of the relevant linear dynamic T within the unit circle in the complex plane. Figure 1a gives an example. This is appropriate for discrete-domain (e.g., discrete-time or discrete-space) dynamics. For continuous-time dynamics, the coronal spectrogram unwraps into what we call the coronated horizon, via the familiar discreteto-continuous conformal mapping of the inside of the unit circle of the complex plane to the left half of the complex plane [16]. Figure 1b displays a discrete-time version of the coronated horizon. Ultimately, either the coronal spectrogram or coronated horizon yield the same information and lend the same important lesson: the eigenvalues of the hidden linear dynamic control allowed system behaviors.

Coronal spectrograms demonstrate that complex systems behave according to the spectrum of their hidden linear dynamic. The relevant frequency-dependent measure $f(\omega)$ emanates from the nonzero eigenvalues of the hidden linear dynamic: the closer eigenvalues approach the unit circle, the sharper the observed peaks. At one extreme, one observes Bragg-like reflections (deltafunction contributions) when the eigenvalues fall on the unit circle. The collection of diffuse peaks observed is a sum of Lorentzian-like and, what we might call, super-Lorentzian-like line profiles. Indeed, the Lorentzian-like line profiles are the discrete-time version of a Lorentzian curve. While the continuous-domain Lorentzian is given by $\operatorname{Re}\left(\frac{c}{\omega-\lambda}\right)$, the continuous-to-discrete conformal mapping $\omega \to e^{i\omega}$ directly yields our discrete-domain analog $\operatorname{Re}\left(\frac{c}{e^{i\omega}-\lambda}\right)$. The super-Lorentzian-like line profiles, from nondiagonalizable contributions, have the form $\operatorname{Re}\left[\left(\frac{c}{e^{i\omega}-\lambda}\right)^n\right].$





(a) (4-3)-GM Process of the (R-k)-Golden (b) (4-3)-GP-(3) Process of the Mean family with $0 \leq k = \nu_0(\zeta) \leq R = \nu_0(\mathcal{W}) < \infty$, which $0 \le k = \nu_0(\zeta) \le \nu_0(\mathcal{W}) < R = \infty$ generates processes with finite but tunable whenever P > 1, generates processes with whenever Z > 1. Markov order is infinite Markov-order R and cryptic-order k. infinite Markov-order R, tunable finite cryptic-order k, and tunable finite symmetry-collapse index $\nu_0(\mathcal{W})$.

(c) (4-3)-GPZ-(3-3) Process of the $(\nu_0(\mathcal{W})-k)$ -Golden Parity-(P) family with $(\nu_0(\mathcal{W})-\nu_0(\zeta))$ -Golden Parity-(P-Z) family with $0 \le \nu_0(\zeta) \le \nu_0(\mathcal{W}) < k = R = \infty$ whenever either P > 1 or Z > 1. Cryptic-order is infinite when Z > 1. This family generates processes with finite but tunable symmetry-collapse index $\nu_0(\mathcal{W})$ and cryptic index $\nu_0(\zeta)$.

FIG. 2. Process families for exploring the roles of and interplay between Markov-order R, cryptic-order k, the symmetrycollapse index $\nu_0(\mathcal{W})$ of the zero eigenvalue of the synchronizing dynamic over mixed states, and the cryptic index $\nu_0(\zeta)$ of the zero eigenvalue of the cryptic operator presentation. We always have $k \leq R$ and $\nu_0(\zeta) \leq \nu_0(\mathcal{W})$. Whenever $\Lambda_{\mathcal{W}} = \Lambda_T \cup \{0\}$, R is finite, $R = \nu_0(\mathcal{W})$ and $k = \nu_0(\zeta)$. Whenever $\Lambda_{\zeta} = \Lambda_T \cup \{0\}$, k is finite, whether or not R is, and $k = \nu_0(\zeta)$. When k or R is infinite, the cryptic index and symmetry-collapse index reveal more nuanced features of the cryptic and synchronization dynamics.

Zero eigenvalues also contribute to $f(\omega)$, but only sinusoidal contributions of discrete increments from $\cos(\omega)$ up to $\cos(\nu_0 \omega)$. Since these are qualitatively distinct from the super-Lorentzian contributions and do not emanate radially from the eigenvalues the same way contributions from nonzero eigenvalues do, coronal spectrograms are most useful for understanding the contributions of nonzero eigenvalues. Nevertheless, the two contributions can be usefully disentangled, as shown later.

We use both coronal spectrograms and coronated horizons to visualize various features in the examples to follow.

VI. **EXAMPLES**

Α. **Golden Mean Processes**

To explore finite Markov order in relation to various complexity measures let's consider the (R-k)-Golden Mean (GM) Processes [17]. This process family describes a unique transition-parametrized process for each Markov-order $R \in \{\nu \in \mathbb{Z} : \nu \geq 1\}$ and each crypticorder $k \in \{\kappa \in \mathbb{Z} : 1 \leq \kappa \leq R\}$. The ϵ -machine for the (4-3)-Golden Mean Process is shown in Fig. 2a. From this the construction of all other (R-k)-Golden Mean processes can be discerned. In words, (R-k)-Golden Mean Processes are binary with alphabet $\mathcal{A} = \{0, 1\}$ and if the most recent history consists of at least k consecutive 0s (and no 1s since then) then there is a probability p of next observing a 1 and a probability 1 - p of simply seeing another 0. The first possibility (observing a 1) entails R consecutive 1s followed by at least k consecutive 0s.

The eigenvalues of the internal state-to-state transition matrix of the ϵ -machine's recurrent component are:

$$\Lambda_T = \left\{ \lambda \in \mathbb{C} : \left(\lambda - (1-p) \right) \lambda^{R+k-1} = p \right\} \,.$$

In the limit of $p \to 1$, all (*R-k*)-Golden Mean Processes become perfectly periodic. In this limit, the eigenvalues are evenly distributed on the unit circle:

$$\Lambda_T \to \left\{ e^{in2\pi/(R+k)} \right\}_{n=0}^{R+k-1}$$

At the other extreme, as $p \to 0$, all eigenvalues evolve to zero, except the stationary eigenvalue at z = 1. At any



FIG. 3. Evolution of eigenvalues Λ_T of the recurrent component of the (5–3)-GM Process's ϵ -machine. Displayed within the unit circle of the complex plane, the trajectory of each eigenvalue follows a line that starts thick blue and ends thin red as the transition parameter p evolves from 1 to 0. In addition to the seven eigenvalues that move from the nontrivial eighth roots of unity towards zero along nonlinear trajectories, the eigenvalue at z = 1 does not change with p.

setting of p, the nonunity eigenvalues lie approximately on a circle within the complex plane whose radius decreases nonlinearly from 1 to 0 as p is swept from 1 to 0. Simultaneously, this circle's center moves from the origin to a positive real value and back to the origin as p is swept from 1 to 0. Figure 3 shows how the eigenvalues of the (5-3)-Golden Mean Process evolve over the full range of p as it sweeps from 1 to 0.

In contrast to the *p*-dependent spectrum of the recurrent structure just discussed, the only eigenvalue corresponding to the transient structure of the S-MSP is equal to zero, regardless of the transition parameter *p*. Recall that this is necessarily true for any process with finite Markov order. Hence, $\Lambda_{\mathcal{W}} = \Lambda_T \cup \{0\}$, with $\nu_0(\mathcal{W}) = R$. The cryptic structure is similar: $\Lambda_{\zeta} = \Lambda_T \cup \{0\}$, with $\nu_0(\zeta) = k$, where ζ is the state-to-state transition matrix of the cryptic operator presentation.

Table I compares the ϵ -machines, autocorrelation, power spectra, MSPs, myopic entropy rates, and myopic state uncertainties for three *p*-parametrized examples of (*R-k*)-GM processes.

The autocorrelation of each process captures their 'leaky periodic' behaviors: The leakiness originates from the self-transition at state A that adds a phase-slip noise to otherwise (R + k)-periodic behavior. Moreover, each process' phase, and so its ϵ -machine's internal state, is uniquely identified after R observations. This corresponds to the depth of the S-MSP tree-like structure $\nu_0(W) = R$, the convergence of the myopic entropy rate $h_{\mu}(L)$ to the true entropy rate h_{μ} when conditioning on observations of finite block-length L - 1 = R, and the complete loss of causal state uncertainty $\mathcal{H}(L)$ after L = R observations.

A paradigm of finite Markov order, the (5-3)-Golden Mean Process has a strictly tree-like structure in its MSP's transients, which have a maximum depth equal to both $\nu(W)$ and its Markov order of 5.

These analyses illustrate the typical behaviors of complexity measures for finite Markov-order processes. We next investigate examples of infinite Markov-order processes to draw attention to the characteristic differences of nonzero eigenvalues in their MSP transient structures.

B. Even Process

The Even Process, shown in the first column of Table II, is a well known example of a stochastic process that cannot be generated by any finite Markov-order approximation, yet it is generated by a simple two-state HMM.

Infinite Markov order, in this case, stems from the fact that the process generates only an even number of consecutive 1s, between 0s. The countably infinite set of Markov chain states necessary to track this parity reflects the infinite order. Moreover, the surplus entropy rate $h_{\mu}(L) - h_{\mu}$ incurred when using a finite order-(L-1)Markov approximation vanishes only asymptotically, being the sum of decaying exponentials. (See Table II.) Such long-lived decay is driven by nonzero eigenvalues in the S-MSP transient structure.

This is in stark contrast to the myopic entropy rate for the finite Markov order processes of Table I. For them $h_{\mu}(L)$ drops to h_{μ} exactly at L = R + 1. Similarly, the average state uncertainty $\mathcal{H}(L)$ for infinite Markov processes converges only asymptotically—and with the same set of decay rates as $h_{\mu}(L)$ —to its asymptotic value of 0. (This curve is not shown in Table II for lack of space.)

The Even Process is a relatively simple example of an infinite Markov-order process. As expected for infinite Markov-order, its MSP's transient structure had nonzero eigenvalues. Generally, though, two ranges of contribution are to be expected in synchronization dynamics. The first is a finite-horizon contribution to the past-future mutual information, corresponding to completely ephemeral zero eigenvalues in the MSP's transient structure. The second is an infinite-horizon contribution to the past-future mutual information, arising from nonzero eigencontributions.



TABLE I. Complexity analyses for finite Markov-order processes. Quantitative data used p = 1/2.

C. Golden–Parity Process Family

To further explore the nature of infinite Markov order processes, we introduce the $(\nu_0 - k)$ -Golden-Parity-(P)Processes. This family subsumes and extends the examples analyzed so far. The role of each parameter is explained in Fig. 2b, which displays a state-transition diagram of the (4-3)-GP-(3) Process' ϵ -machine. If P = 1, the family reduces to the $(\nu_0 - k)$ -Golden Mean Process family, with tunable Markov $R = \nu_0(W)$ and cryptic k orders. That is, $(\nu_0(W)-k)$ -GP- $(1) = (\nu_0(W)-k)$ -GM. However, the Markov order becomes infinite whenever P > 1. In this case the index $\nu_0(W)$ of the S-MSP's zero-eigenvalue—which controls the finite duration necessary to resolve all broken symmetries and the cryptic order k can still be tuned independently.



TABLE II. Complexity analyses for infinite Markov-order processes. Quantitative data used p = 1/2 and q = 1/3.

The Even Process considered earlier is the (0-0)-Golden-Parity-(2) Process.

Three examples of $(\nu_0(\mathcal{W})\text{-}k)\text{-}\text{Golden-Parity-}(P)$ processes are analyzed in Table II. The S-MSP transient structure for the second two clarifies the difference between (i) the symmetry collapse associated with completely ephemeral transient states that are fully depleted

of probability density after $\nu_0(\mathcal{W})$ time steps and (ii) the long-lived leaky transients whose probability density only vanishes as more-refined ambiguity is resolved.

Examining the myopic entropy convergence $h_{\mu}(L)$, the effect of these distinct routes to synchronization on the predictability can be seen: The process is much more predictable, on average, after $\nu_0(\mathcal{W})$ time steps. How-

ever, the average predictability of an infinite-Markovorder process continues to increase with increasing observation window, albeit with exponentially diminishing returns. In general, we showed that this asymptotic convergence occurs as a sum of decaying exponentials from diagonalizable subspaces and as the product of polynomials and exponentials in the case of nondiagonalizable structures associated with nonzero eigenvalues. The apparent oscillations under the exponential decays are completely described by the leaky periodicities of the eigenvalues in the transient belief states.

Finally, note that the excess entropy spectrum $\mathcal{E}(\omega)$ shows the frequency domain view of observation-induced predictability. $\mathbf{E} = \lim_{\omega \to 0} \mathcal{E}(\omega)$ is the total past-future mutual information, which is also the excess entropy observed before full synchronization. The $\nu_0(\mathcal{W})$ symmetry collapse contributes significantly and early to the total excess entropy of the last two examples. Whereas, the asymptotic tails of synchronization associated with leaky periodicity of particular transient states of uncertainty accumulate their contribution to excess entropy rather slowly.

In addition to new intuitions about convergence behaviors in stochastic processes, the general and broadly applicable theoretical results here allow novel numerical investigations and unprecedentedly-accurate analyses of infinite-Markov-order processes. As an example of the latter, let us summarize several of the exact results derived in App. A for the (p, q)-parametrized (2-1)-GP-(2) process explored in Table II's second column.

Depending on whether the transition parameter p is larger or smaller than $2\sqrt{q} - q$, App. A found qualitatively distinct behaviors dominate the (2-1)-GP-(2) process. This hints at a general principle: behaviorally distinct regions are separated by a critical line in the (p,q)parameter space along which the transition dynamic Tbecomes nondiagonalizable. For $p < 2\sqrt{q} - q$, the autocorrelation for $|L| \geq 2$ has the exact solution:

$$\gamma(L) = \beta^2 + \beta q^{|L|/2} \operatorname{Re}(\zeta e^{i\omega_{\xi}|L|}) , \qquad (37)$$

where $\beta \equiv 2(p+2q)/(1+p+2q), \ \zeta \equiv (\xi+1)^2(p\xi+2q)/(\xi(\xi^3+p\xi+2q)), \ \xi \equiv -\frac{1}{2}(p+q)+i\frac{1}{2}\sqrt{4q-(p+q)^2},$ and $\omega_{\xi} \equiv \frac{\pi}{2} + \arctan((p+q)/\sqrt{4q-(p+q)^2})$. The corresponding power spectrum is:

$$P(\omega) = \frac{8q}{1+p+2q} + \frac{2p}{1+p+2q} [1 - \cos(\omega)] + \beta \operatorname{Re}\left(\frac{\zeta\xi}{e^{i\omega} - \xi} + \frac{\zeta\xi}{e^{-i\omega} - \xi}\right) + 2\pi\beta^2 \sum_{k=-\infty}^{\infty} \delta(\omega + 2\pi k) .$$
(38)

For any parameter setting, the metadynamic of observation-induced synchronization to the (2-1)-GP-(2) process is nondiagonalizable due to the index-2 zero eigenvalue. This leads to a completely ephemeral contribution to $h_{\mu}(L)$ up to L = 2. For $L \geq 3$, we find the myopic entropy rate relaxes asymptotically to the true entropy rate according to:

$$h_{\mu}(L) - h_{\mu} = \begin{cases} \frac{-p \log p + (1+p) \log(1+p) - 2p}{\sqrt{p}(1+p+2q)} p^{L/2} & \text{for odd } L, \\ \frac{p \log p - (1+p) \log(1+p) + 2}{(1+p+2q)} p^{L/2} & \text{for even } L, \end{cases}$$

where the process' true entropy rate is:

$$h_{\mu} = \frac{-q\log q - p\log p - (1 - p - q)\log(1 - p - q)}{1 + p + 2q}$$

Interestingly, while the autocorrelation at separation L scales as $\sim q^{L/2}$, the predictability of transitions between single-symbol-shifted histories of length L converges as $\sim p^{L/2}$ —indicating two rather independent decay rates.

The amount of the future that can be predicted from the past is the total mutual information between the observable past and observable future:

$$\mathbf{E} = \frac{(1-p-q)\log(1-p-q)-p\log p-q\log q-(1-p)\log(1-p)}{1+p+2q} + \log(1+p+2q) \ .$$

However, to actually *perform* prediction requires more memory than this amount of shared information. Calculation of additional measures and more detail can be found in App. A.

To explore the structure in infinite-cryptic order processes, one can use the more generalized family of $(\nu_0(W)-\nu_0(\zeta))$ -GP-(P-Z) Processes. For them, $\nu_0(\zeta)$ is the index of the zero-eigenvalue of the cryptic operator presentation and the process has infinite cryptic order whenever Z > 1. Above, Z = 1 and $(\nu_0(W)-\nu_0(\zeta))$ -GP- $(P-1) = (\nu_0(W)-\nu_0(\zeta))$ -GP-(P). Since the preceding examples served well enough to illustrate the power of spectral decomposition, our main goal, we leave a full analysis of this family to interested others.

VII. PREDICTING SUPERPAIRWISE STRUCTURE

The Random-Random-XOR (RRXOR) Process is generated by a simple HMM. Figure 4 displays its fivestate ϵ -machine. However, it illustrates nontrivial, counterintuitive features typical of stochastic dynamic information processing systems. The process is defined over three steps that repeat: (i) a 0 or 1 is output with equal probability, (ii) another 0 or 1 is output with equal prob-



FIG. 4. RRXOR Process ϵ -machine.

ability, and then (iii) the eXclusive-OR operation (XOR) of the last two outputs is output.

Surprisingly, but calculations easily verify, there are no pairwise correlations. All of its correlations are higher than second order. One consequence is that its power spectrum is completely flat—the signature of white noise; see Fig. 5. This would lead a casual observer to incorrectly conclude that the generated time series has no structure. In fact, a white noise spectrum is an indication that, if structure it present, it must be hidden in higher-order correlations.

The RRXOR Process clearly is not structureless—via the exclusive OR, it transforms information in a substantial way. We show that the complexity measures introduced above can detect this higher-order structure. However, let us first briefly consider *why* the correlationbased measures fail to detect structure in the RRXOR Process.

It is sometimes noted that information measures are superior to standard measures of correlation since they capture nonlinear dependencies, while the standard correlation relies on linear models. And so, we can avoid this problem by using the *information correlation* $I[X_0; X_\tau]$ rather than autocorrelation. Analogous to autocorrelation, it too has a spectral version—the *power-of-pairwise information (POPI) spectrum*:

$$\mathcal{I}(\omega) \equiv -\mathrm{H}(X_0) + \lim_{N \to \infty} \sum_{\tau = -N}^{N} e^{-i\omega\tau} \mathrm{I}[X_0; X_{\tau}] . \quad (39)$$

It is easy to show that $\mathcal{I}(\omega) = 0$ for the RRXOR Process. Hence, as Fig. 5 showed, such measures are still not sufficient to detect even simple computational structure, since they only can detect *pairwise* statistical dependencies.

In stark contrast, the excess entropy spectrum $\mathcal{E}(\omega)$ does identify the structure of hidden dependencies in the RRXOR Process; see Fig. 6. Why? The brief de-



FIG. 5. Power spectrum $P(\omega)$ and POPI spectrum $\mathcal{I}(\omega)$ of the RRXOR Process: The first is flat and the second identically zero. One might incorrectly conclude the RRXOR Process is structureless white noise.

tour through power spectra, information correlation, and POPI spectra brings us to a deeper understanding of why $\mathcal{E}(\omega)$ is successful at detecting nuanced computational structure in a time series. Since it partitions all random variables throughout time, the excess entropy itself picks up any systematic influence the past has on the future. The excess entropy *spectrum* further identifies the frequency decomposition of any such linear or nonlinear dependencies. In short, all multivariate dependencies contribute to the excess entropy spectrum.

Let us now consider the hidden structure of the RRXOR Process in more detail. With reference to Fig. 4, we observe that the expected probability density over causal states evolves through the ϵ -machine with a period-3 modulation. In a given realization, the particular symbols emitted after each phase resetting ($\phi = 0$) break symmetries with respect to which "wings" of the ϵ -machine structure are traversed. This is reflected in T's eigenvalues: the three roots of unity $\{e^{in2\pi/3}\}_{n=0}^2$ and two zero eigenvalues, with $a_0(T) = g_0(T) = 2$ giving index $\nu_0(T) = 1$.

The period-3 modulation leads to a phase ambiguity when an observer synchronizes to the process—an ambiguity resolved in the MSP transient structure. This resolution is rather complicated, as made explicit in the RRXOR Process' S-MSP, shown in Fig. 7. There are 31 transient states of uncertainty, in addition to the five recurrent states—36 mixed states in total.

Since we derived the ϵ -machine's S-MSP, W = W. Hence, the MSP's layout depicts the information processing involved while an observer synchronizes to the RRXOR Process. This graphically demonstrates the burden on an optimal predictor—the prerequisite paths to synchronization before an observer is surprised by the



FIG. 6. Excess entropy spectrum of the RRXOR Process, together with the eigenvalues of the S-MSP transition matrix W. Among the power spectrum, POPI spectrum, and excess entropy spectrum, only the excess entropy spectrum is able to detect structure in the RRXOR Process since the structure is beyond pairwise. The eigenspectrum of the MSP of the RRXOR ϵ -machine and the excess entropy spectrum both indicate that the RRXOR Process is indeed structured, with both ephemeral symmetry breaking and leaky periodicities in the convergence to optimal predictability.

irreducible uncertainty of only h_{μ} bits per observation, and before the observer only needs to heed the relevant b_{μ} bits per observation to stay synchronized, on average.

The MSP introduces new, relevant zero eigenvalues associated with its transient states. In particular, the firstencountered tree-like transients (starting with mixedstate π) introduce new Jordan blocks up to dimension 2. Overall, the 0-eigenspace of W has index 2, so that $\nu_0(W) = 2$.

Two different sets of leaky-period-3 structures appear in the MSP transients. There are four leaky three-state cycles, each with the same leaky-period-3 contributions to the spectrum: $\left\{ \left(\frac{1}{4}\right)^{1/3} e^{in2\pi/3} \right\}_{n=0}^2$. There are also four leaky four-state cycles, each with a leaky-period-3 contribution and symmetry-breaking 0-eigenvalue contribution to the spectrum: $\left\{ \left(\frac{1}{2}\right)^{1/3} e^{in2\pi/3} \right\}_{n=0}^2 \cup \{0\}$. The difference in eigenvalue magnitude, $\left(\frac{1}{4}\right)^{1/3}$ versus $\left(\frac{1}{2}\right)^{1/3}$, implies different timescales of synchronization associated with distinct learning tasks. For example, an immediate lesson is that it takes longer (on average) to escape the



FIG. 7. MSP of the RRXOR Process' ϵ -machine: Grayed out (and dashed) transitions permanently leave the states from which they came. Recognizing the manner by which these transitions partition the mixed-state space allows simplified spectrum calculations. The directed graph structure is inherently nonplanar. The large blue recurrent state should be visualized as being *behind* the transient states; it does not contain them.

4-state leaky-period-3 components (from the time of arrival) than to escape the preceding 3-state leaky-period-3 components of the synchronizing metadynamic.

The entropy rate convergence plots of Figs. 8 and 9 reveal a sophisticated predictability modulation that simply could not have been gleaned from the spectra of Fig. 5. Figure 9 emphasizes the dominance of the slowest-decaying eigenmodes for large L. Such oscillations under the exponential convergence to synchronization are typical. However, as seen in comparison with Fig. 8 much of the uncertainty may be reduced before this asymptotic mode comes to dominate. Ultimately, synchronization to optimal prediction may involve important contributions from all modes of the mixed-state-to-state metadynamic.

This detailed analysis of the RRXOR Process suggests several general lessons about how we view information in stochastic processes. First, as information processing increases in sophistication, a vanishing amount of a process' intrinsic structure will be discernible at low-orders of correlation. Second, logical computation, as implemented by universal logic gates, primarily operates above pairwise correlation. And so, finally, there is substantial motivation to move beyond measures of pairwise corre-



FIG. 8. Ephemeral $(h_{\multimap}(L))$ and persistent $(h_{\leadsto}(L))$ contributions to the myopic entropy rate $(h_{\mu}(L))$. The ephemeral contribution lasts only up to $L = \nu_0(W) = 2$.



FIG. 9. Tails of the myopic entropy convergence $h_{\mu}(L)$ shown in Fig. 8 decay according to two different leaky period-three envelopes. The latter correspond to the two qualitatively different types of transient synchronization cycles in the MSP of Fig. 7. One of the transient cycles has a relatively fast decay rate of $r_2 = (1/4)^{1/3}$. While the slower decay rate of $r_1 = (1/2)^{1/3}$ dominates $h_{\mu}(L)$'s deviation from h_{μ} at large L.

VIII. CONCLUSION

Surprisingly, many questions we ask about structured stochastic, nonlinear processes implicate a linear dynamic over an appropriate hidden state space. That is, there is an implied hidden Markov model. The promise is that once the dynamic is found for the question of interest, one can make progress in analyzing it. Unfortunately, a roadblock immediately arises: these hidden linear dynamics are generically nondiagonalizable for questions related to prediction and to information and complexity measures. Deploying Part I's meromorphic functional calculus, though, circumvents the roadblock. Using it, we determined closed-form expressions for a very wide range of information and complexity measures. Often, these expressions turned out to be direct functions of the HMM's transition dynamic.

This allowed us to catalog in detail the range of possible convergence behaviors for correlation and myopic uncertainty. The analytic formulas revealed a new symmetry-collapse index that is a lower bound on the Markov order and serves as an independent, more nuanced timescale of information processing. We then considered complexity measures that accumulate during the transient relaxation to observer synchronization. We also introduced the new notion of complexity spectra, gave a new kind of information-theoretic signal analysis in terms of coronal spectrograms, and highlighted common simplifications for special cases, such as almost diagonalizable dynamics. We closed by analyzing several families of finite and infinite Markov and cryptic order processes and emphasized the importance of higher-than-pairwise-order correlations, showing how the excess entropy spectrum is the key diagnostic tool for them.

The analytical completeness might suggest that we have reached an end. Partly, but the truth we seek is rather farther down the road. The meromorphic functional calculus of nondiagonalizable operators merely sets the stage for the next challenges—to develop complexity measures and structural decompositions for infinite-state and infinite excess entropy processes. Hopefully, the new toolset will help us scale the hierarchies of truly complex processes outlined in Refs. [1, 11, 12], at a minimum giving exact answers at each stage of a convergent series of finite- ϵ -machine approximations.

ACKNOWLEDGMENTS

lation. We must learn to recognize hidden structures and to use higher-order structural investigations to better understand information processing. This is critical to empirically probing functionality in biological and engineered processes.

The authors thank Alec Boyd, Chris Ellison, Ryan James, John Mahoney, and Dowman Varn for helpful discussions. JPC thanks the Santa Fe Institute for its hospitality. This material is based upon work supported by,



FIG. 10. ϵ -Machine of the (2-1)-GP-(2) Process.

or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under contract numbers W911NF-12-1-0234, W911NF-13-1-0340, and W911NF-13-1-0390.

Appendix A: Example Analytical Calculations

To exercise the operational nature of the framework introduced, the following explicitly carries out the analytic calculations to obtain the closed-form complexity measures for the (p, q)-parametrized (2-1)-GP-(2) Process. This process was already visually explored in the second column of Table II. And so, the goal here is primarily pedagogical—providing insight and better explicating particular calculational steps. The appendix demonstrates a variety of techniques in the spirit of a tutorial, though many were not called out in the main development.

1. Process and spectra features

The (p,q)-parametrized (2-1)-GP-(2) Process is described by its ϵ -machine, whose state-transition diagram was shown in the first row and second column of Table II and is reproduced here in Fig. 10. Formally, the (2-1)-GP-(2) stationary stochastic process is generated by the HMM $\mathcal{M}_{\epsilon M} = (\mathcal{S}, \mathcal{A}, \{T^{(x)}\}_{x \in \mathcal{A}}, \eta_0 = \pi)$. That is, \mathcal{M} consists of a set of hidden causal states $\mathcal{S} = \{A, B, C, D\}$, an alphabet $\mathcal{A} = \{0, 1, 2\}$ of symbols emitted to form the observed process, and a set $\{T^{(x)}: T^{(x)}_{s,s'} = \Pr(X_t = x, \mathcal{S}_{t+1} = s' | \mathcal{S}_t = s)\}_{x \in \mathcal{A}}$ of symbol-labeled transition matrices. These are:

and

$$T^{(2)} = \begin{bmatrix} 0 & 0 & q & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

The symbol-labeled transition matrices sum to the rowstochastic internal state-to-state transition matrix:

$$T = \sum_{x \in \mathcal{A}} T^{(x)} = \begin{bmatrix} 1 - p - q & p & q & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix} .$$

This is a Markov chain over the hidden states. From it we find the stationary state distribution $\langle \pi | = \langle \pi | T$:

$$\langle \pi | = \frac{1}{1+p+2q} \begin{bmatrix} 1 & p & q \end{bmatrix} .$$

The first analysis task is to determine the eigenvalues and associated projection operators for the internal stateto-state transition matrix T. From:

$$det(T - \lambda I) = \lambda(\lambda - 1) \left(\lambda^2 + (p + q)\lambda + q\right)$$
$$= 0,$$

we find the four eigenvalues:

$$\Lambda_T = \left\{1, \, 0, \, -\frac{1}{2}(p+q) \pm \frac{1}{2}\sqrt{(p+q)^2 - 4q} \right\} \,.$$

All eigenvalues are real for $p \ge 2\sqrt{q} - q$. Two are complex with nonzero imaginary part when $p < 2\sqrt{q} - q$. Putting this together with the transition-probability consistency constraint that p + q < 1 yields the map of the transition-parameter space shown in Fig. 11.

For a generic choice of the parameter setting (p,q), all *T*'s eigenvalues are unique, and so *T* is diagonalizable. However, two of the eigenvalues become degenerate along the parameter subspace $p = 2\sqrt{q} - q$, giving $\Lambda_T = \{1, 0, -\sqrt{q}\}$. We find that the algebraic multiplicity $a_{-\sqrt{q}} = 2$ is larger than the geometric multiplic-



FIG. 11. Transition-probability space for the (p,q)-parametrized (2-1)-GP-(2) Process.

ity $g_{-\sqrt{q}} = 1$, yielding nondiagonalizability ($\nu_{-\sqrt{q}} = 2$) along this ($p = 2\sqrt{q} - q$)-submanifold. More broadly, experience has shown that eigenvalues generically induce nondiagonalizability when they collide and scatter in the complex plane. For example, this occurs when a pair of eigenvalues first "entangle" to become complex conjugate pairs.

For any parameter setting (p,q), we find T's right eigenvectors from $(T - \lambda I) |\lambda\rangle = \vec{0}$. They are:

$$\left|0\right\rangle = \begin{bmatrix} 0 & 1 & -p/q & 0 \end{bmatrix}^{\top}$$

and

$$|\lambda\rangle = \begin{bmatrix} \lambda^2 & \lambda & 1 & \lambda \end{bmatrix}^{\top}$$

for all $\lambda \in \Lambda_T \setminus \{0\}$. Similarly, we find the left eigenvectors of T from $\langle \lambda | (T - \lambda I) = \vec{0}$:

$$\langle 0| = \begin{bmatrix} 0 & 1 & 0 & -1 \end{bmatrix}$$

and

$$\langle \lambda | = \begin{bmatrix} \lambda^2 & p\lambda & q\lambda & q \end{bmatrix} ,$$

for all $\lambda \in \Lambda_T \setminus \{0\}$. Clearly, the left eigenvectors are *not* simply the complex-conjugate transpose of the right eigenvectors. This is a signature of the more intricate algebraic structure in these processes.

Since T is generically diagonalizable, all of the projection operators are simply the normalized ket-bra outer

products:

$$T_{\lambda} = \frac{|\lambda\rangle \langle \lambda|}{\langle \lambda | \lambda \rangle} ,$$

so long as $p \neq 2\sqrt{q} - q$. To wit:

$$T_1 = \frac{|1\rangle \langle 1|}{\langle 1|1\rangle}$$
$$= |1\rangle \langle \pi|$$

Along the nondiagonalizable $(p = 2\sqrt{q} - q)$ -subspace of parameter settings: $\langle -\sqrt{q} | -\sqrt{q} \rangle = 0$. This corresponds to the fact that the right and left eigenvectors are now dual to the left and right *generalized* eigenvectors, rather than being dual to each other. We find the projection operator for the nondiagonalizable eigenspace is:

$$T_{-\sqrt{q}} = \frac{|-\sqrt{q}^{(1)}\rangle \langle -\sqrt{q}^{(2)}|}{\langle -\sqrt{q}^{(2)}| -\sqrt{q}^{(1)}\rangle} + \frac{|-\sqrt{q}^{(2)}\rangle \langle -\sqrt{q}^{(1)}|}{\langle -\sqrt{q}^{(1)}| -\sqrt{q}^{(2)}\rangle} ,$$

where:

$$\begin{split} |{}^{-}\sqrt{q}{}^{(1)}\rangle &= \begin{bmatrix} q & -\sqrt{q} & 1 & -\sqrt{q} \end{bmatrix}^{\top} , \\ |{}^{-}\sqrt{q}{}^{(2)}\rangle &= \begin{bmatrix} -\sqrt{q} & 0 & 1/\sqrt{q} & 0 \end{bmatrix}^{\top} , \\ \langle{}^{-}\sqrt{q}{}^{(1)}| &= \begin{bmatrix} q & q(\sqrt{q}-2) & -q\sqrt{q} & q \end{bmatrix} , \text{ and} \\ \langle{}^{-}\sqrt{q}{}^{(2)}| &= \begin{bmatrix} -\sqrt{q} & 0 & 0 & \sqrt{q} \end{bmatrix} . \end{split}$$

Equivalently, $T_{-\sqrt{q}} = I - T_1 - T_0$. The projection operators for the remaining $(\lambda = 0, 1)$ eigenspaces retain the same form as before: $T_{\lambda} = |\lambda\rangle \langle \lambda| / \langle \lambda|\lambda\rangle$.

2. Observed correlation

The pieces are now in place to calculate the observable correlation and power spectrum. Recall that we derived the general spectral decomposition of the autocorrelation function $\gamma(L) = \langle \overline{X}_t X_{t+L} \rangle_t$:

$$\gamma(L) = \sum_{\substack{\lambda \in \Lambda_T \\ \lambda \neq 0}} \sum_{m=0}^{\nu_{\lambda}-1} \langle \pi \overline{\mathcal{A}} | T_{\lambda,m} | \mathcal{A} \mathbf{1} \rangle \binom{|L|-1}{m} \lambda^{|L|-1-m} + [0 \in \Lambda_T] \sum_{m=0}^{\nu_0-1} \langle \pi \overline{\mathcal{A}} | T_0 T^m | \mathcal{A} \mathbf{1} \rangle \, \delta_{|L|-1,m}$$

for nonzero integer L, where:

$$\begin{aligned} \langle \pi \overline{\mathcal{A}} | &= \langle \pi | \sum_{x \in \mathcal{A}} \overline{x} T^{(x)} \\ &= \frac{1}{1 + p + 2q} \begin{bmatrix} p & p & 2q & 2q \end{bmatrix} \end{aligned}$$

,

and:

$$\begin{aligned} |\mathcal{A}\mathbf{1}\rangle &= \sum_{x \in \mathcal{A}} x T^{(x)} |\mathbf{1}\rangle \\ &= \begin{bmatrix} p + 2q \ 1 \ 2 \ 0 \end{bmatrix}^{\top} \end{aligned}$$

For generic parameter settings, this reduces to:

$$\begin{split} \gamma(L) &= \frac{\langle \pi \overline{\mathcal{A}} | \, 0 \rangle \langle 0 \, | \mathcal{A} \mathbf{1} \rangle}{\langle 0 | 0 \rangle} \delta_{|L|,1} \\ &+ \sum_{\lambda \in \Lambda_T \atop \lambda \neq 0} \frac{\langle \pi \overline{\mathcal{A}} | \, \lambda \rangle \langle \lambda \, | \mathcal{A} \mathbf{1} \rangle}{\langle \lambda | \lambda \rangle} \lambda^{|L|-1} \\ &= 4 \left(\frac{p+2q}{1+p+2q} \right)^2 + \frac{-p}{1+p+2q} \delta_{|L|,1} \\ &+ \frac{p+2q}{1+p+2q} \sum_{\lambda \in \Lambda_T \atop \lambda \neq 0,1} \frac{(\lambda+1)^2(p\lambda+2q)}{\lambda^3 + p\lambda + 2q} \lambda^{|L|-1} \;. \end{split}$$

Notably, the autocorrelation splits into an ephemeral part (via the Kronecker delta) due to the zero eigenvalue, an exponentially decaying oscillatory part due to the two eigenvalues with magnitude between 0 and 1, and an asymptotic part that survives even as $L \to \infty$ due to the eigenvalue of unity. Moreover, for $p < 2\sqrt{q} - q$, the two nontrivial eigenvalues become complex conjugate pairs. This allows us to rewrite the autocorrelation for the (2-1)-GP-(2) process for $|L| \geq 2$ concisely as:

$$\gamma(L) = \beta^2 + \beta q^{|L|/2} \operatorname{Re}(\zeta e^{i\omega_{\xi}|L|}) ,$$

where:

$$\begin{split} \beta &\equiv \frac{2(p+2q)}{1+p+2q} \ , \\ \zeta &\equiv \frac{(\xi+1)^2(p\xi+2q)}{\xi(\xi^3+p\xi+2q)} \ , \\ \xi &\equiv -\frac{1}{2}(p+q) + i\frac{1}{2}\sqrt{4q-(p+q)^2} \ , \text{ and} \\ \omega_\xi &\equiv \frac{\pi}{2} + \arctan\left(\frac{p+q}{\sqrt{4q-(p+q)^2}}\right) \ . \end{split}$$

The latter form reveals that the magnitude of the largest nonunity eigenvalue $|\xi| = \sqrt{q}$ controls the slowest rate of decay of observed correlation $|\gamma(L)| \sim \sqrt{q}^L$ and that the complex phase of the eigenvalues determine the oscillations within this exponentially decaying envelope.

We found the general spectral decomposition of the continuous part of the power spectrum is:

$$P_{\rm c}(\omega) = \left\langle |x|^2 \right\rangle + 2 \sum_{\lambda \in \Lambda_T} \sum_{m=0}^{\nu_{\lambda}-1} \operatorname{Re} \frac{\left\langle \pi \overline{\mathcal{A}} | T_{\lambda,m} | \mathcal{A} \mathbf{1} \right\rangle}{(e^{i\omega} - \lambda)^{m+1}} .$$

Also, recall from earlier that the δ -functions of the power spectrum arise from T's eigenvalues that lie on the unit circle:

$$P_{\rm d}(\omega) = \sum_{k=-\infty}^{\infty} \sum_{\substack{\lambda \in \Lambda_T \\ |\lambda|=1}} 2\pi \,\delta(\omega - \omega_{\lambda} + 2\pi k) \\ \times \operatorname{Re}\left(\lambda^{-1} \langle \pi \overline{\mathcal{A}} | T_{\lambda} | \mathcal{A} \mathbf{1} \rangle\right)$$

where ω_{λ} is related to λ by $\lambda = e^{i\omega_{\lambda}}$. As long as p+q < 1, $\lambda = 1$ is the only eigenvalue that lies on the unit circle, so that:

$$P_{\rm d}(\omega) = \sum_{k=-\infty}^{\infty} 2\pi \,\delta(\omega + 2\pi k) \operatorname{Re} \langle \pi \overline{\mathcal{A}} | T_1 | \mathcal{A} \mathbf{1} \rangle \ ,$$

with:

Re
$$\langle \pi \overline{\mathcal{A}} | T_1 | \mathcal{A} \mathbf{1} \rangle = \beta^2$$

= $4 \left(\frac{p + 2q}{1 + p + 2q} \right)^2$

Putting this all together we find the complete power spectrum for the (2-1)-GP-(2) Process:

$$\begin{split} P(\omega) &= \frac{8q}{1+p+2q} + \frac{2p}{1+p+2q} \big[1 - \cos(\omega) \big] \\ &+ \beta \sum_{\substack{\lambda \in \Lambda_T \\ \lambda \neq 0, 1}} \operatorname{Re} \left(\frac{(\lambda+1)^2 (p\lambda+2q)}{\lambda^3 + p\lambda + 2q} \cdot \frac{1}{e^{i\omega} - \lambda} \right) \\ &+ 2\pi \beta^2 \sum_{k=-\infty}^{\infty} \delta(\omega + 2\pi k) \;, \end{split}$$

Thus, the process exhibits δ -functions from the eigenvalue on the unit circle, continuous Lorentzian-like line profiles emanating from finite eigenvalues inside the unit circle which express the process' chaotic nature, and the unique sinusoidal contribution that can only come from zero eigenvalues.

Whenever $p < 2\sqrt{q} - q$, the two nontrivial eigenvalues become complex conjugate pairs. This allows us to rewrite the process' power spectrum in a more transparent way:

$$P(\omega) = \frac{8q}{1+p+2q} + \frac{2p}{1+p+2q} \left[1 - \cos(\omega)\right] + \beta \operatorname{Re}\left(\frac{\zeta\xi}{e^{i\omega} - \xi} + \frac{\zeta\xi}{e^{-i\omega} - \xi}\right) + 2\pi\beta^2 \sum_{k=-\infty}^{\infty} \delta(\omega + 2\pi k) ,$$

which is clearly symmetric about $\omega = 0$.

The level of completeness achieved is notable: we cal-

culated these properties exactly in closed form for this infinite-Markov order stochastic process over the full range of possible parameter settings. Moreover, once it constructs the process' S-MSP, the next section goes on to produce the same level of analytical completeness but for predictability.

3. Predictability

To analyze the process' predictability, we need the S-MSP of any of its generators. Since we started already with the ϵ -machine, we directly determine its S-MSP. This gives the mixed-state transition matrix W = W that, in turn, suffices for calculating both predictability in this section and the synchronization necessary for prediction in the next.

We construct the S-MSP by calculating all mixed states that can be induced by observation from the start distribution π and then calculating the transition probabilities between them. There are eight such mixed-state distributions over \boldsymbol{S} . However, four of them (δ_A , δ_B , δ_C , and δ_D) correspond to completely synchronized peaked distributions. The other four states are new (relative to the recurrent states) transient states and correspond to transient states of recurrent-state uncertainty during synchronization. Calculating, we find the eight unique mixed-state distributions iteratively from:

$$\left\langle \eta^{wx} \right| = \frac{\left\langle \eta^w \right| T^{(x)}}{\left\langle \eta^w \right| T^{(x)} \left| \mathbf{1} \right\rangle} \ ,$$

starting with:

$$\langle \eta^x | = \frac{\langle \pi | T^{(x)}}{\langle \pi | T^{(x)} | \mathbf{1} \rangle} ,$$

are:

$$\begin{split} \langle \pi | &= \frac{1}{1+p+2q} \begin{bmatrix} 1 & p & q & q \end{bmatrix} ,\\ \langle \eta^1 | &= \frac{1}{2} \begin{bmatrix} 1 & 1 & 0 & 0 \end{bmatrix} ,\\ \langle \eta^{11} | &= \frac{1}{1+p} \begin{bmatrix} 1 & p & 0 & 0 \end{bmatrix} ,\\ \langle \eta^2 | &= \frac{1}{2} \begin{bmatrix} 0 & 0 & 1 & 1 \end{bmatrix} ,\\ \langle \delta_A | &= \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix} ,\\ \langle \delta_B | &= \begin{bmatrix} 0 & 1 & 0 & 0 \end{bmatrix} ,\\ \langle \delta_C | &= \begin{bmatrix} 0 & 0 & 1 & 0 \end{bmatrix} , \text{ and }\\ \langle \delta_D | &= \begin{bmatrix} 0 & 0 & 0 & 1 \end{bmatrix} . \end{split}$$

In this, the transient mixed states η^w are labeled according to the shortest word w that induces them. These



FIG. 12. S-MSP of the (2-1)-GP-(2) Process.

distributions constitute the set of mixed-states \mathcal{R}_{π} of the S-MSP. Moreover, each transition probability from mixed state $\langle \eta^w | to mixed state \langle \eta^{wx} |$ is calculated as: $\langle \eta^w | T^{(x)} | \mathbf{1} \rangle$. Altogether, these calculations yield the S-MSP of the (2-1)-GP-(2) Process, reproduced in Fig. 12 from Table II for convenience.

As a HMM, the ϵ -machine's S-MSP is specified by the 4-tuple: $\mathcal{M}_{\text{S-MSP}} = (\mathcal{R}_{\pi}, \mathcal{A}, \{\mathcal{W}^{(x)}\}_{x \in \mathcal{A}}, \mu_0 = \delta_{\pi}),$ where \mathcal{R}_{π} is the set of mixed states just quoted, \mathcal{A} is the same observable alphabet as before, $\{\mathcal{W}^{(x)}\}_{x \in \mathcal{A}}$ is the set of symbol-labeled transition matrices among the mixed states, and $\langle \delta_{\pi} | = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$ is the start distribution over the mixed states.

With this new linear metadynamic in hand, our next step is to calculate the eigenvalues and projection operators of the internal mixed-state-to-state transition dynamic $\mathcal{W} = \sum_{x \in \mathcal{A}} \mathcal{W}^{(x)}$. \mathcal{W} can be explicitly represented in the block-matrix form:

$$\mathcal{W} = \begin{bmatrix} A & B \\ \mathbf{0} & T \end{bmatrix}$$

where:

$$A = \begin{bmatrix} 0 & \frac{2p}{1+p+2q} & 0 & \frac{2q}{1+p+2q} \\ 0 & 0 & \frac{1+p}{2} & 0 \\ 0 & \frac{2p}{1+p} & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \text{ and}$$
$$B = \begin{bmatrix} \frac{1-p}{1+p+2q} & 0 & 0 & 0 \\ \frac{1-p-q}{2} & 0 & \frac{q}{2} & 0 \\ \frac{1-p-q}{1+p} & 0 & \frac{q}{1+p} & 0 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{bmatrix},$$

and T is the same as the state-to-state internal transition matrix of the ϵ -machine from earlier.

 \mathcal{W} 's eigenvalues are thus relatively straightforward to

calculate since $\det(\mathcal{W} - \lambda I) = \det(A - \lambda I)\det(T - \lambda I)$ implies that:

$$\Lambda_{\mathcal{W}} = \Lambda_A \cup \Lambda_T \; .$$

The new eigenvalues introduced by the feedback matrix A are $\Lambda_A = \{0, \pm \sqrt{p}\}$ with $\nu_0 = 2$. The new nonzero eigenvalues $(\pm \sqrt{p})$, found most easily from Part I's cyclic eigenvalue rule, are associated with diagonalizable subspaces. It is important to note that, while T was only nondiagonalizable along a very special submanifold in parameter space, the mixed-state-to-state metadynamic is generically nondiagonalizable over all parameter settings. This nondiagonalizability corresponds to a special kind of symmetry breaking of uncertainty during synchronization.

W's eigenvectors are most easily found through a twostep process. Specifically, $|{}^{\pm}\sqrt{p}_{A}\rangle$ and $\langle{}^{\pm}\sqrt{p}_{A}|$ (the solutions of $A |{}^{\pm}\sqrt{p}_{A}\rangle = \pm \sqrt{p} |{}^{\pm}\sqrt{p}_{A}\rangle$ and $\langle{}^{\pm}\sqrt{p}_{A}|A = \pm \sqrt{p} \langle{}^{\pm}\sqrt{p}_{A}|)$ are found first, and the result is used to reduce the number of unknowns when solving the full eigenequations $(\mathcal{W} |{}^{\pm}\sqrt{p}_{\mathcal{W}}\rangle = \pm \sqrt{p} |{}^{\pm}\sqrt{p}_{\mathcal{W}}\rangle)$ for $|{}^{\pm}\sqrt{p}_{\mathcal{W}}\rangle$ and $\langle{}^{\pm}\sqrt{p}_{\mathcal{W}}|$. Similarly, we can recycle the restricted eigenvectors $|\lambda_{T}\rangle$ and $\langle\lambda_{T}|$ found earlier for the ϵ -machine to reduce the number of unknowns when solving the more general eigenvector problems for $|\lambda_{\mathcal{W}}\rangle$ and $\langle\lambda_{\mathcal{W}}|$ in cases where $\lambda \in \Lambda_{T}$. Performing such a calculation, we find:

$$\begin{split} |^{\pm}\sqrt{p}_{\mathcal{W}}\rangle &= \begin{bmatrix} \frac{1}{1+p+2q} & \frac{\pm\sqrt{p}}{2p} & \frac{1}{1+p} & 0 & 0 & 0 & 0 \end{bmatrix}^{\top} ,\\ \langle^{\pm}\sqrt{p}_{\mathcal{W}}| &= \begin{bmatrix} 0 & \pm\sqrt{p} & \frac{1+p}{2} & 0 & -\frac{1\pm\sqrt{p}}{2} & -\frac{p\pm\sqrt{p}}{2} & 0 & 0 \end{bmatrix} ,\\ \langle\lambda_{\mathcal{W}}| &= \begin{bmatrix} 0 & 0 & 0 & \lambda^2 & p\lambda & q\lambda & q \end{bmatrix} \text{ for all } \lambda \in \Lambda_T \setminus \{0\} ,\\ |1_{\mathcal{W}}\rangle &= \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}^{\top} , \text{ and} \\ |\lambda_{\mathcal{W}}\rangle &= \begin{bmatrix} 0 & \frac{\lambda(1+\lambda)}{2} & \frac{\lambda(1+\lambda)}{1+p} & \frac{1+\lambda}{2} & \lambda^2 & \lambda & 1 & \lambda \end{bmatrix}^{\top} , \end{split}$$

for $\lambda \in \Lambda_T \setminus \{0, 1\}$. Moreover, \mathcal{W} 's eigenvectors and generalized eigenvector corresponding to eigenvalue 0 are:

$$\begin{split} |0_1^{(1)}\rangle &= \langle \delta_{\pi}|^{\top} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}^{\top} ,\\ |0_1^{(2)}\rangle &= \begin{bmatrix} 0 & 0 & \frac{1+p+2q}{2q} & 0 & 0 & 0 \end{bmatrix}^{\top} ,\\ |0_2^{(1)}\rangle &= \begin{bmatrix} 0 & \frac{1}{2} & \frac{p}{1+p} & \frac{-p}{2q} & 0 & 1 & \frac{-p}{q} & 0 \end{bmatrix}^{\top} ,\\ \langle 0_1^{(1)}| &= \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 & \frac{-1}{2} & \frac{-1}{2} \end{bmatrix} ,\\ \langle 0_1^{(2)}| &= \begin{bmatrix} \frac{1+p+2q}{2q} & 0 & \frac{-(1+p)}{2q} & 0 & 0 & 0 & \frac{-1}{2} & \frac{-1}{2} \end{bmatrix} , \text{ and }\\ \langle 0_2^{(1)}| &= \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & -1 \end{bmatrix} . \end{split}$$

Above, we used the notation $|0_k^{(m)}\rangle$ for indexing generalized eigenvectors introduced in Part I.

All nondegenerate eigenvalues have projection operators of the form:

$$\mathcal{W}_{\lambda} = rac{\ket{\lambda}ra{\lambda}}{ra{\lambda}ra{\lambda}}$$

However, the degenerate and nondiagonalizable subspace associated with the zero eigenvalue has the composite projection operator:

$$\mathcal{W}_0 = \frac{|0_1^{(1)}\rangle \langle 0_1^{(2)}|}{\langle 0_1^{(2)}|0_1^{(1)}\rangle} + \frac{|0_1^{(2)}\rangle \langle 0_1^{(1)}|}{\langle 0_1^{(1)}|0_1^{(2)}\rangle} + \frac{|0_2^{(1)}\rangle \langle 0_2^{(1)}|}{\langle 0_2^{(1)}|0_2^{(1)}\rangle} \ .$$

The fact that $\langle \delta_{\pi} | \lambda_{\mathcal{W}} \rangle = 0$ for all $\lambda \in \Lambda_T \setminus (\{1\} \cup \Lambda_A)$ is an instantiation of a general result that greatly simplifies the calculations relating to predictability and prediction.

The remaining piece for analyzing predictability is the vector of transition-entropies $|\mathcal{H}(W^{\mathcal{A}})\rangle$. A simple calculation, utilizing the fact that:

$$-\sum_{i} \frac{n_i}{d} \log\left(\frac{n_i}{d}\right) = \log(d) - \frac{1}{d} \sum_{i} n_i \log n_i$$

when $\sum_{i} n_i = d$, yields:

$$|\mathcal{H}(W^{\mathcal{A}})\rangle = \begin{bmatrix} \log(1+p+2q) - \frac{1}{1+p+2q} \left[2q\log(2q) + 2p\log(2p) + (1-p)\log(1-p) \right] \\ 1 - \frac{1}{2} \left[q\log(q) + (1+p)\log(1+p) + (1-p-q)\log(1-p-q) \right] \\ \log(1+p) - \frac{1}{1+p} \left[q\log(q) + 2p\log(2p) + (1-p-q)\log(1-p-q) \right] \\ 1 \\ - \left[q\log(q) + p\log(p) + (1-p-q)\log(1-p-q) \right] \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

where log is understood to be the base-2 logarithm \log_2 .

Putting this all together, we can now calculate in full detail the myopic entropy rate $h_{\mu}(L)$ that results from modeling the infinite-order (2-1)-GP-(2) process as an order-(L-1) Markov process:

$$\begin{split} h_{\mu}(L) &= \sum_{\lambda \in \Lambda_{W} \atop \lambda \neq 0} \sum_{m=0}^{\nu_{\lambda}-1} \left\langle \delta_{\pi} \right| W_{\lambda,m} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \binom{L-1}{m} \lambda^{L-1-m} \\ &+ \left[0 \in \Lambda_{W} \right] \sum_{m=0}^{\nu_{0}-1} \delta_{L-1,m} \left\langle \delta_{\pi} \right| W_{0} W^{m} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \\ &= \delta_{L,1} \left\langle \delta_{\pi} \right| W_{0} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle + \delta_{L,2} \left\langle \delta_{\pi} \right| WW_{0} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \\ &+ \left\langle \delta_{\pi} \right| W_{1} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle + \sum_{\lambda \in \Lambda_{W} \setminus \Lambda_{T}} \left\langle \delta_{\pi} \right| W_{\lambda} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \lambda^{L-1}. \end{split}$$

After L = 2, this teeters along an exponentially decaying envelope as it approaches its asymptotic value of:

$$h_{\mu} = \langle \delta_{\pi} | W_{1} | H(W^{\mathcal{A}}) \rangle$$

= $\langle \delta_{\pi} | 1_{\mathcal{W}} \rangle \langle \pi_{\mathcal{W}} | H(W^{\mathcal{A}}) \rangle$
= $\langle \pi_{\mathcal{W}} | H(W^{\mathcal{A}}) \rangle = \langle \pi_{T} | H(T^{\mathcal{A}}) \rangle$
= $\frac{-q \log q - p \log p - (1 - p - q) \log(1 - p - q)}{1 + p + 2q}$

Simplifying the terms in the myopic entropy rate yields:

$$\begin{split} \left\langle \delta_{\pi} \right| W_{0} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle &= \left\langle \delta_{\pi} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \\ &- \frac{1+p}{1+p+2q} \left\langle \delta_{\eta^{11}} \left| \mathbf{H}(W^{\mathcal{A}}) \right\rangle \end{split}$$

and:

$$\langle \delta_{\pi} | WW_0 | \mathbf{H}(W^{\mathcal{A}}) \rangle = \frac{2q}{1+p+2q}$$

for the two ephemeral contributions.

For $L \geq 3$, we find for odd L:

$$h_{\mu}(L) - h_{\mu} = \frac{-p\log p + (1+p)\log(1+p) - 2p}{\sqrt{p}(1+p+2q)} p^{L/2}$$

and for even L:

$$h_{\mu}(L) - h_{\mu} = \frac{p \log p - (1+p) \log(1+p) + 2}{1+p+2q} p^{L/2}$$
.

This highlights the period-2 nature of the asymptotic decay.

The total mutual information between the observable

past and observable future, the excess entropy, is:

$$\mathbf{E} = \sum_{\lambda \in \Lambda_W \setminus \{1\}} \sum_{m=0}^{\nu_{\lambda}-1} \frac{1}{(1-\lambda)^{m+1}} \langle \delta_{\pi} | W_{\lambda,m} | \mathbf{H}(W^{\mathcal{A}}) \rangle$$

= $\frac{(1-p-q)\log(1-p-q) - p\log p}{1+p+2q}$
 $-\frac{q\log q + (1-p)\log(1-p)}{1+p+2q}$
 $+\log(1+p+2q)$.

This is the total future information that can possibly be predicted using past observations. The structure of *how* this information is unraveled over time is revealed in the excess entropy spectrum from $-\pi < \omega \leq \pi$:

$$\begin{split} \mathcal{E}(\omega) &= \pi h_{\mu} \delta(\omega) \\ &+ \sum_{m=0}^{\nu_0 - 1} \cos\left((m+1)\,\omega\right) \left\langle \delta_{\pi} | \, W_0 W^m \, | \mathbf{H}(W^{\mathcal{A}}) \right\rangle \\ &+ \sum_{\lambda \in \Lambda_W \setminus 0} \sum_{m=0}^{\nu_{\lambda} - 1} \operatorname{Re} \left(\frac{\left\langle \delta_{\pi} | \, W_{\lambda,m} \, | \mathbf{H}(W^{\mathcal{A}}) \right\rangle}{(e^{i\omega} - \lambda)^{m+1}} \right) \\ &= \pi h_{\mu} \delta(\omega) + \frac{2q}{1+p+2q} \cos(2\omega) \\ &+ \left(\mathbf{E} + \frac{p \log p - (1+p) \log(1+p) - 2q}{1+p+2q} \right) \cos(\omega) \\ &+ \frac{-p \log p + (1+p) \log(1+p)}{1+p+2q} \operatorname{Re} \left(\frac{e^{i\omega} - p}{e^{i2\omega} - p} \right) \\ &+ \frac{2p}{1+p+2q} \operatorname{Re} \left(\frac{1-e^{i\omega}}{e^{i2\omega} - p} \right) \,. \end{split}$$

From this, we observe that $\mathbf{E} = \lim_{\omega \to 0} \mathcal{E}(\omega)$.

4. Synchronizing to predict optimally

To analyze the information-processing cost of synchronizing to a process, we need its ϵ -machine S-MSP. We already constructed this in the last section. Hence, we can immediately evaluate the resources necessary for synchronizing to and predicting the (2-1)-GP-(2) Process.

The only novel piece needed for this is the vector of mixed-state entropies $|H[\eta]\rangle$. A simple calculation yields:

$$|\mathbf{H}[\eta]\rangle = \begin{bmatrix} \log(1+p+2q) - \frac{p\log p + 2q\log q}{1+p+2q} \\ 1 \\ \log(1+p) - \frac{p\log p}{1+p} \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

,

where log is again understood to be the base-2 logarithm $\log_2.$

An observer, tasked with predicting the future as well as possible, must synchronize to the causal state of the dynamic. The remaining causal-state uncertainty $\mathcal{H}^+(L)$ after an observation interval of L steps is:

$$\begin{aligned} \mathcal{H}^{+}(L) \\ &= \sum_{\lambda \in \Lambda_{\mathcal{W}} \atop \lambda \neq 0} \sum_{m=0}^{\nu_{\lambda}-1} \left\langle \delta_{\pi} | \mathcal{W}_{\lambda,m} | \mathbf{H}[\eta] \right\rangle \binom{L}{m} \lambda^{L-m} \\ &+ \left[0 \in \Lambda_{\mathcal{W}} \right] \sum_{m=0}^{\nu_{0}-1} \delta_{L,m} \left\langle \delta_{\pi} | \mathcal{W}_{0} \mathcal{W}^{m} | \mathbf{H}[\eta] \right\rangle \\ &= \delta_{L,0} \left[\log(1+p+2q) - \frac{2q \log q + (1+p) \log(1+p)}{1+p+2q} \right] \\ &+ \delta_{L,1} \frac{2q}{1+p+2q} + \frac{\left\langle \sqrt{p} | \mathbf{H}[\eta] \right\rangle + (-1)^{L} \left\langle -\sqrt{p} | \mathbf{H}[\eta] \right\rangle}{1+p+2q} p^{L/2} . \end{aligned}$$

More explicitly, for $L \ge 2$, this becomes:

$$\mathcal{H}^{+}(L) = \begin{cases} \frac{2\sqrt{p}}{1+p+2q} p^{L/2} & \text{for odd } L\\ \frac{(1+p)\log(1+p)-p\log p}{1+p+2q} p^{L/2} & \text{for even } L \end{cases}.$$

The total synchronization information accumulated is then:

$$\mathbf{S} = \sum_{\lambda \in \Lambda_{\mathcal{W}} \setminus \{1\}} \sum_{m=0}^{\nu_{\lambda}-1} \frac{1}{(1-\lambda)^{m+1}} \left\langle \delta_{\pi} | \mathcal{W}_{\lambda,m} | \mathbf{H}[\eta] \right\rangle$$
$$= \frac{2q(1-\log q)}{1+p+2q} + \frac{p[2-\log p+(1+p)\log(1+p)]}{(1-p)(1+p+2q)}$$
$$+ \log(1+p+2q) .$$

Even after synchronization, an observer must update an average of b_{μ} of its bits of information per observation and must keep track of a net C_{μ} bits of information to stay synchronized, where:

$$C_{\mu} = \mathbf{H}[\pi]$$

= $\langle \delta_{\pi} | \mathbf{H}[\eta] \rangle$
= $\log(1 + p + 2q) - \frac{p \log p + 2q \log q}{1 + p + 2q}$.

An interesting feature of prediction is a process' crypticity:

$$\chi = C_{\mu} - \mathbf{E}$$

= $\frac{3q \log q + (1-p) \log(1-p)}{1+p+2q}$
 $-\frac{(1-p-q) \log(1-p-q)}{1+p+2q}$

This gives, as a function of p and q, the minimal overhead of additional memory about the past—beyond the information that the future can possibly share with it—that must be stored for optimal prediction.

In summary, we now more fully appreciate, via this rather complete analysis, the fundamental limits on predictability of our example stochastic process. It showed many of the qualitative features, both in terms of the calculation and system behavior, that should be expected when analyzing prediction, based on the more general results of the main development. The procedures can be commandeered to apply to any inference algorithm that yields a generative model, whether classical machine learning, Bayesian Structural Inference [9], or any other favorite inference tool. With a generative model in hand, synchronizing to real world data—necessary to make good predictions about the real world's futurefollows the S-MSP metadynamics. The consequence for prediction is that typically there will be a finite epoch of symmetry collapse followed by a slower asymptotic synchronization that allows improved prediction, as longer observations induce a refined knowledge of what lies hidden.

- S. E. Marzen and J. P. Crutchfield. Nearly maximally predictive features and their dimensions. *Phys. Rev. E*, 95(5):051301(R), 2017. 2, 21
- [2] P. M. Riechers and J. P. Crutchfield. Beyond the spectral theorem: Decomposing arbitrary functions of nondiagonalizable operators. arxiv.org:1607.06526. 2, 9, 10
- [3] C. E. Shannon and W. Weaver. The Mathematical

Theory of Communication. University of Illinois Press, Champaign-Urbana, 1962.

- [4] A. V. Oppenheim and R. W. Schafer. Discrete-Time Signal Processing. Prentice Hall Press, Upper Saddle River, NJ, USA, 3rd edition, 2009.
- [5] D.P. Varn and J.P. Crutchfield. Chaotic crystallography: how the physics of information reveals structural order

in materials. *Current Opinion in Chemical Engineering*, 7:47 – 56, 2015. Biological engineering / Materials engineering.

- [6] Ya. G. Sinai. On the notion of entropy of a dynamical system. Dokl. Akad. Nauk. SSSR, 124:768, 1959.
- [7] J. P. Crutchfield and N. H. Packard. Symbolic dynamics of noisy chaos. *Physica*, 7D(1-3):201 – 223, 1983.
- [8] S. G. Williams. Introduction to symbolic dynamics. In Proceedings of Symposia in Applied Mathematics, volume 60, pages 1–12, 2004. 2, 29
- [9] C. C. Strelioff and J. P. Crutchfield. Bayesian structural inference for hidden processes. *Phys. Rev. E*, 89:042119, Apr 2014. 2, 28
- [10] T. M. Cover and M. E. Hellman. The two-armed-bandit problem with time-invariant finite memory. *IEEE Trans. Info. Th.*, 16(2):185–195, Mar 1970. 6
- [11] J. P. Crutchfield and S. E. Marzen. Signatures of infinity: Nonergodicity and resource scaling in prediction, complexity, and learning. *Phys. Rev. E*, 91:050106(R), 2015. 6, 21
- [12] J. P. Crutchfield and D. P. Feldman. Regularities unseen, randomness observed: Levels of entropy convergence. CHAOS, 13(1):25–54, 2003. 7, 10, 12, 21
- [13] J. P. Crutchfield, C. J. Ellison, J. R. Mahoney, and R. G. James. Synchronization and control in intrinsic and de-

- [14] P. M. Riechers, D. P. Varn, and J. P. Crutchfield. Diffraction patterns of layered close-packed structures from hidden markov models. arxiv.org:1410.5028. 11, 12
- [15] D. Lind and B. Marcus. An Introduction to Symbolic Dynamics and Coding. Cambridge University Press, New York, 1995. 12
- [16] B. Latni. Signal processing and linear systems. Oxford University Press, USA, 1998. 13
- [17] Historically, the name "Golden Mean" derives from the (1-1)-Golden Mean Process that forbids two consecutive 1s. It describes the symbolic dynamics of the chaotic onedimensional shift map over the unit interval with generating partition, when the slope of the map is the golden mean $\phi = 1+1/\phi = (1+\sqrt{5})/2$ [8]. Moreover, consistency with the stationary stochastic process generated implies $1 - p = 1/\phi$. Generalizing to arbitrary p and especially to arbitrary R and k takes us away from the eponymous property into a setting sufficiently general to study the generic behavior of order-R Markov processes and the complexities of predicting them. 14